

# Robust Group Linkage

Pei Li  
University of Zurich  
peili@ifi.uzh.ch

Xin Luna Dong \*  
Google Inc.  
lunadong@google.com

Songtao Guo \*  
LinkedIn Corp.  
songtao.gg@gmail.com

Andrea Maurino  
University of Milan-Bicocca  
maurino@disco.unimib.it

Divesh Srivastava  
AT&T Labs-Research  
divesh@research.att.com

## ABSTRACT

We study the problem of *group linkage*: linking records that refer to entities in the same group. Applications for group linkage include finding businesses in the same chain, finding conference attendees from the same affiliation, finding players from the same team, etc. Group linkage faces challenges not present for traditional record linkage. First, although different members in the same group can share some similar *global* values of an attribute, they represent different entities so can also have distinct *local* values for the same or different attributes, requiring a high *tolerance* for value diversity. Second, groups can be huge (with tens of thousands of records), requiring high *scalability* even after using good blocking strategies.

We present a two-stage algorithm: the first stage identifies *cores* containing records that are very likely to belong to the same group, while being robust to possible erroneous values; the second stage collects strong evidence from the cores and leverages it for merging more records into the same group, while being tolerant to differences in local values of an attribute. Experimental results show the high effectiveness and efficiency of our algorithm on various real-world data sets.

## 1. INTRODUCTION

*Record linkage* aims at linking records that refer to the same real-world entity and it has been extensively studied in the past years (surveyed in [8, 19]). In this paper we study a related but different problem that we call *group linkage*: linking records that refer to entities in the same group.

One major motivation for our work comes from identifying *business chains*—connected business entities that share a brand name and provide similar products and services (e.g., *Walmart*, *McDonald's*). With the advent of the Web and mobile devices, we are observing a boom in *local search*; that is, searching local businesses under geographical constraints. Local search engines include *Google Maps*, *Yahoo! Local*, *YellowPages*, *yelp*, *ezlocal*, etc. The knowledge of business chains can have a big economic value to local search engines, as it allows users to search by business chain, allows search engines to render the returned results by chains, allows data collectors to clean and enrich information within the same chain, allows the associated review system to connect reviews

\*Research conducted at AT&T Labs-Research.

**Table 1: Identified top-5 US business chains. For each chain, we show the number of stores, distinct business names, distinct phone numbers, distinct URL domain names, and distinct categories.**

Name	#Store	#Name	#Phn	#URL	#Cat
<i>SUBWAY</i>	21,912	772	21,483	6	23
<i>Bank of America</i>	21,727	48	6,573	186	24
<i>U-Haul</i>	21,638	2,340	18,384	14	20
<i>USPS - United State Post Office</i>	19,225	12,345	5,761	282	22
<i>McDonald's</i>	17,289	2401	16,607	568	47

on branches of the same chain, and allows sales people to target potential customers. Business listings are rarely associated with specific chains explicitly in real-world business-listing collections, so we need to identify the chains. Sharing the same name, phone number, or URL domain name can all serve as evidence of belonging to the same chain. However, for US businesses alone there are tens of thousands of chains and as we show soon, we cannot easily develop any rule set that applies to all chains.

We are also motivated by applications where we need to find people from the same organization, such as counting conference attendees from the same affiliation, counting papers by authors from the same institution, and finding players of the same team. The organization information is often missing, incomplete, or simply too heterogeneous to be recognized as the same (e.g., “International Business Machines Corporation”, “IBM Corp.”, “IBM”, “IBM Research Labs”, “IBM-Almaden”, etc., all refer to the same organization). Contact phones, email addresses, and mailing addresses of people all provide extra evidence for group linkage, but they can also vary for different people even in the same organization.

Group linkage faces challenges not present for traditional record linkage. First, although different members in the same group can share some similar *global* values of an attribute, they represent different entities so can also have distinct *local* values for the same or different attributes. For example, different branches in the same business chain can provide different local phone numbers, different addresses, etc. It is non-trivial to distinguish such differences from various representations for the same value and sometimes erroneous values in the data. Second, there are often millions of records for group linkage, and a group can contain tens of thousands of members. A good blocking strategy should put these tens of thousands of records in the same block; but performing record linkage via traditional pairwise comparisons within such huge blocks can be very expensive. Thus, *scalability* is a big challenge. We use the following example of identifying business chains throughout the paper for illustration.

**EXAMPLE 1.1.** *We consider a set of 18M real-world business listings in the US extracted from Yellowpages.com, each describing a business by its name, phone number, URL domain name, location, and category. Our algorithm automatically finds 600K business chains and 2.7M listings that belong to these chains. Table 1 lists*

**Table 2: Real-world business listings. We show only state for location and simplify names of category. There is a wrong value in *italic font*.**

RID	name	phone	URL (domain)	location	category
$r_1$	Home Depot, The	808		NJ	furniture
$r_2$	Home Depot, The	808		NY	furniture
$r_3$	Home Depot, The	808	homedepot	MD	furniture
$r_4$	Home Depot, The	808	homedepot	AK	furniture
$r_5$	Home Depot, The	808	homedepot	MI	furniture
$r_6$	Home Depot, The	101	homedepot	IN	furniture
$r_7$	Home Depot, The	102	homedepot	NY	furniture
$r_8$	Home Depot, USA	103	homedepot	WV	furniture
$r_9$	Home Depot USA	808		SD	furniture
$r_{10}$	Home Depot - Tools	808		FL	furniture
$r_{11}$	Taco Casa		tacocasa	AL	restaurant
$r_{12}$	Taco Casa	900	tacocasa	AL	restaurant
$r_{13}$	Taco Casa	900	tacocasa, <i>tacocasatexas</i>	AL	restaurant
$r_{14}$	Taco Casa	900		AL	restaurant
$r_{15}$	Taco Casa	900		AL	restaurant
$r_{16}$	Taco Casa	701	tacocasatexas	TX	restaurant
$r_{17}$	Taco Casa	702	tacocasatexas	TX	restaurant
$r_{18}$	Taco Casa	703	tacocasatexas	TX	restaurant
$r_{19}$	Taco Casa	704		NY	food store
$r_{20}$	Taco Casa		tacodelmar	AK	restaurant

the largest five chains we found. We observe that (1) each chain contains up to 22K different branch stores, (2) different branches from the same chain can have a large variety of names, phone numbers, and URL domain names, and (3) even chains of similar sizes can have very different numbers of distinct URLs (same for other attributes). Thus, rule-based linkage can hardly succeed and scalability is a necessity.

Table 2 shows a set of 20 business listings (with some abstraction) in this data set. After investigating their webpages manually, we find that  $r_1 - r_{18}$  belong to three business chains:  $Ch_1 = \{r_1 - r_{10}\}$ ,  $Ch_2 = \{r_{11} - r_{15}\}$ , and  $Ch_3 = \{r_{16} - r_{18}\}$ ;  $r_{19}$  and  $r_{20}$  do not belong to any chain. Note the slightly different names for businesses in chain  $Ch_1$ ; also note that  $r_{13}$  is integrated from different sources and contains two URLs, one (tacocasatexas) being wrong.

Simple linkage rules do not work well on this data set. For example, if we require only high similarity on name for chain identification, we may wrongly decide that  $r_{11} - r_{20}$  all belong to the same chain as they share a popular restaurant name Taco Casa. Traditional linkage strategies do not work well either. If we apply Swoosh-style linkage [27] and iteratively merge records with high similarity on name and shared phone or URL, we can wrongly merge  $Ch_2$  and  $Ch_3$  because of the wrong URL from  $r_{13}$ . If we require high similarity between listings on name, phone, URL, category, we may either split  $r_6 - r_8$  out of chain  $Ch_1$  because of their different local phone numbers, or learn a low weight for phone but split  $r_9 - r_{10}$  out of chain  $Ch_1$  since sharing the same phone number, the major evidence, is downweighted.  $\square$

The key idea in our solution is to find strong evidence that can glue group members together, while being tolerant to differences in values specific for individual group members. For example, we wish to reward sharing of primary values, such as *primary phone numbers* or *URL domain names* for chain identification, but would not penalize differences from local values, such as *locations*, local phone numbers, and even categories. For this purpose, our algorithm proceeds in two stages. First, we identify *cores* containing records that are very likely to belong to the same group. Second, we collect strong evidence from the resulting cores, such as primary phone numbers and URL domain names in business chains, based on which we cluster the cores and remaining records into groups. The use of cores and strong evidence distinguishes our clustering algorithm from traditional clustering techniques for record linkage. In this process, it is crucial that core generation makes very few

false positives even in the presence of erroneous values, such that we can avoid ripple effect on clustering later. Our algorithm is designed to ensure efficiency and scalability.

The group linkage problem we study in this paper is different from the group linkage in [18, 23], which decides similarity between *pre-specified* groups of records. Our goal is to find records that belong to the same group and we make three contributions.

1. We study core generation in presence of erroneous data. Our core is *robust* in the sense that even if we remove a few possibly erroneous records from a core, we still have strong evidence that the rest of the records in the core must belong to the same group.
2. We then reduce the group linkage problem into clustering cores and remaining records. Our clustering algorithm leverages strong evidence collected from cores and meanwhile is *tolerant* to value variety of records in the same group.
3. We conducted experiments on two real-world data sets in different domains, showing high efficiency and effectiveness of our algorithms.

Note that we assume prior to group linkage, we first conduct record linkage (e.g., [15]). Our experiments show that minor mistakes for record linkage do not significantly affect the results of group linkage, and records that describe the same entity but fail to be merged in the record-linkage step are often put into the same group. We plan to study how to combine record linkage and group linkage to improve the results of both in the future.

In the rest of the paper, Section 2 discusses related work. Section 3 defines the problem and provides an overview of our solution. Sections 4-5 describe the two stages in our solution. Section 6 describes experimental results. Section 7 concludes.

## 2. RELATED WORK

Record linkage has been extensively studied in the past (surveyed in [8, 19]). Traditional linkage techniques aim at linking records that refer to the same real-world entity, so implicitly assume value consistency between records that should be linked. Group linkage is different in that it aims at linking records that refer to different entities in the same group. The variety of individual entities requires better use of strong evidence and tolerance on different values even within the same group. These two features differentiate our work from any previous linkage technique.

For record clustering in linkage, existing work may apply the transitive rule [17], or do match-and-merge [27], or reduce it to an optimization problem [16]. Our work is different in that our core-identification algorithm aims at being robust to a few erroneous records; and our clustering algorithm emphasizes leveraging the strong evidence collected from the cores.

For record-similarity computation, existing work can be rule based [17], classification based [11], or distance based [6]. There has also been work on weight (or model) learning from labeled data [11, 29]. Our work is different in that in addition to learning a weight for each attribute, we also learn a weight for each value based on whether it serves as important evidence for the group. Note that some previous works are also tolerant to different values but leverage evidence that may not be available in our contexts: [10] is tolerant to schema heterogeneity from different relations by specifying matching rules; [15] is tolerant to possibly false values by considering agreement between different data providers; [21] is tolerant to out-of-date values by considering time stamps; we are tolerant to diversity within the same group.

Two-stage clustering has been proposed in the IR and machine learning community [1, 20, 22, 28, 30]; however, they identify

cores in different ways. Techniques in [20, 28] consider a core as a single record, either randomly selected or selected according to the weighted degrees of nodes in the graph. Techniques in [30] generate cores using agglomerative clustering but can be too conservative and miss strong evidence. Techniques in [1] identify cores as *bi-connected components*, where removing any node would not disconnect the graph. Although this corresponds to the *1-robustness* requirement in our solution (defined in Section 4), they generate overlapping clusters; it is not obvious how to derive non-overlapping clusters in applications such as business-chain identification and how to extend their techniques to guarantee *k-robustness*. Finally, techniques in [20, 22] require knowledge of the number of clusters for one of the stages, so do not directly apply in our context. We compare with these methods whenever applicable in experiments (Section 6), showing that our algorithm is robust in presence of erroneous values and consistently generates high-accuracy results on data sets with different features.

Finally, we distinguish our work from the *group linkage* in [18, 23], which has different goals. On et al. [23] decided similarity between pre-specified groups of records and the group-entity relationship is many-to-many (e.g., authors and papers). Huang [18] decided whether two pre-specified *groups* of records from different data sources refer to the same group by analysis of social network. Our goal is to find records that belong to the same group.

### 3. OVERVIEW

This section formally defines the group linkage problem and provides an overview of our solution.

#### 3.1 Problem definition

Let  $\mathbf{R}$  be a set of records that describe real-world entities by a set of attributes  $\mathbf{A}$ . For each record  $r \in \mathbf{R}$ , we denote by  $r.A$  its value on attribute  $A \in \mathbf{A}$ . Sometimes a record may contain erroneous or missing values.

We consider the *group linkage* problem; that is, finding records that represent entities belonging to the same real-world group. As an example application, we wish to find *business chains*—a set of business entities with the same or highly similar names that provide similar products and services (e.g., *Walmart*, *Home Depot*, *Subway* and *McDonald's*).<sup>1</sup> We focus on non-overlapping groups, which often hold in applications.

**DEFINITION 3.1 (GROUP LINKAGE).** *Given a set  $\mathbf{R}$  of records, group linkage identifies a set of clusters  $\mathbf{CH}$  of records in  $\mathbf{R}$ , such that (1) records that represent real-world entities in the same group belong to one cluster, and (2) records from different groups belong to different clusters.*  $\square$

**EXAMPLE 3.2.** *Consider records in Example 1.1, where each record describes a business store (at a distinct location) by attributes name, phone, URL, location, and category.*

*The ideal solution to the group linkage problem contains 5 clusters:  $Ch_1 = \{r_1 - r_{10}\}$ ,  $Ch_2 = \{r_{11} - r_{15}\}$ ,  $Ch_3 = \{r_{16} - r_{18}\}$ ,  $Ch_4 = \{r_{19}\}$ , and  $Ch_5 = \{r_{20}\}$ . Among them,  $Ch_2$  and  $Ch_3$  represent two different chains with the same name.*  $\square$

#### 3.2 Overview of our solution

Group linkage is related to but different from traditional record linkage because it essentially looks for records that represent entities in the same group, rather than records that represent exactly the same entity. Different members in the same group often share a certain amount of commonality (e.g., common name, primary phone,

and URL domain of chain stores), but meanwhile can also have a lot of differences (e.g., different addresses, local phone numbers, and local URL domains); thus, we need to allow much higher variety in some attribute values to avoid false negatives. On the other hand, as we have shown in Example 1.1, simply lowering our requirement on similarity of records or similarity of a few attributes in clustering can lead to a lot of false positives.

The key intuition of our solution is to distinguish between *strong* evidence and *weak* evidence. For example, different branches in the same business chain often share the same URL domain name and those in North America often share the same 1-800 phone number. Thus, a URL domain or phone number shared among many business listings with highly similar names can serve as strong evidence for chain identification. In contrast, a phone number shared by only a couple of business entities is much weaker evidence, since one might be an erroneous or out-of-date value.

To facilitate leveraging strong evidence, our solution consists of two stages. The first stage collects records that are highly likely to belong to the same group; for example, a set of business listings with the same name and phone number are very likely to be in the same chain. We call the results *cores* of the groups; from them we can collect strong evidence such as name, primary phone number, and primary URL domain of chains. The key goal of this stage is to be robust against erroneous values and make as few false positives as possible, so we can avoid identifying strong evidence wrongly and causing incorrect ripple effect later; however, we need to keep in mind that being too strict can miss important strong evidence.

The second stage clusters cores and remaining records into groups according to the discovered strong evidence. It decides whether several cores belong to the same group, and whether a record that does not belong to any core actually belongs to some group. It also employs weak evidence, but treats it differently from strong evidence. The key intuition of this stage is to leverage the strong evidence and meanwhile be tolerant to diversity of values in the same group, so we can reduce false negatives made in the first stage.

We next illustrate our approach for business-chain identification.

**EXAMPLE 3.3.** *Continue with the motivating example. In the first stage we generate three cores:  $Cr_1 = \{r_1 - r_7\}$ ,  $Cr_2 = \{r_{14}, r_{15}\}$ ,  $Cr_3 = \{r_{16} - r_{18}\}$ . Records  $r_1 - r_7$  are in the same core because they have the same name, five of them ( $r_1 - r_5$ ) share the same phone number 808 and five of them ( $r_3 - r_7$ ) share the same URL homedepot. Similar for the other two cores. Note that  $r_{13}$  does not belong to any core, because one of its URLs is the same as that of  $r_{11} - r_{12}$ , and one is the same as that of  $r_{16} - r_{18}$ , but except name, there is no other common information between these two groups of records. To avoid mistakes, we defer the decision on  $r_{13}$ . Indeed, recall that *tacocasatexas* is a wrong value for  $r_{13}$ . For a similar reason, we defer the decision on  $r_{12}$ .*

*In the second stage, we generate groups—business chains. We merge  $r_8 - r_{10}$  with core  $Cr_1$ , because they have similar names and share either the primary phone number or the primary URL. We also merge  $r_{11} - r_{13}$  with core  $Cr_2$ , because (1)  $r_{12} - r_{13}$  share the primary phone 900 with  $Cr_2$ , and (2)  $r_{11}$  shares the primary URL *tacocasa* with  $r_{12} - r_{13}$ . We do not merge  $Cr_2$  and  $Cr_3$  though, because they share neither the primary phone nor the primary URL. We do not merge  $r_{19}$  or  $r_{20}$  to any core, because there is again not much strong evidence. We thus obtain the ideal result.*  $\square$

To facilitate this two-stage solution, we find attributes that provide evidence for group identification and classify them into three categories.

<sup>1</sup>[http://en.wikipedia.org/wiki/Chain\\_store](http://en.wikipedia.org/wiki/Chain_store).

- *Common-value attribute*: We call an attribute  $A$  a common-value attribute if all entities in the same group have the same or highly similar  $A$ -values. Such attributes include **business-name** for chain identification and **organization** for organization linkage.
- *Dominant-value attribute*: We call an attribute  $A$  a dominant-value attribute if entities in the same group often share one or a few primary  $A$ -values (but there can also exist other less-common values), and these values are seldom used by entities outside the group. Such attributes include **phone** and **URL-domain** for chain identification, and **office-address**, **phone-prefix**, and **email-server** for organization linkage.
- *Multi-value attribute*: We call the rest of the attributes multi-value attributes as there is often a many-to-many relationship between groups and values of these attributes. Such attributes include **category** for chain identification.

The classification can be either learned from training data based on cardinality of attribute values, or performed by domain experts since there are typically only a few such attributes.

We describe core identification in Section 4 and group linkage in Section 5. Our algorithms require common-value and dominant-value attributes, which typically exist for groups in practice. While we present the algorithms for the setting of one machine, a lot of components of our algorithms can be easily parallelized in Hadoop infrastructure [24, 4]; it is not the focus of the paper and we briefly describe the opportunities in Section 6.4.

## 4. CORE IDENTIFICATION

The first stage of our solution creates cores consisting of records that are very likely to belong to the same group. The key goal in core identification is to be robust to possible erroneous values. This section starts with presenting the criteria we wish the cores to meet (Section 4.1), then describes how we efficiently construct similarity graphs to facilitate core finding (Section 4.2), and finally gives the algorithm for core identification (Section 4.3). Note that the notations in this section can be slightly different from those in Graph Theory.

### 4.1 Criteria for a core

At the first stage we wish to make only decisions that are highly likely to be correct; thus, we require that each core contains only highly similar records, and different cores are fairly different and easily distinguishable from each other. In addition, we wish that our results are robust even in the presence of a few erroneous values in the data. In the motivating example,  $r_1 - r_7$  form a good core, because *808* and *homedepot* are very popular values among these records. In contrast,  $r_{13} - r_{18}$  do not form a good core, because records  $r_{14} - r_{15}$  and  $r_{16} - r_{18}$  do not share any phone number or URL domain; the only “connector” between them is  $r_{13}$ , so they can be wrongly merged if  $r_{13}$  contains erroneous values. Also, considering  $r_{13} - r_{15}$  and  $r_{16} - r_{18}$  as two different cores is risky, because (1) it is not very clear whether  $r_{13}$  is in the same chain as  $r_{14} - r_{15}$  or as  $r_{16} - r_{18}$ , and (2) these two cores share one URL domain name so are not fully distinguishable.

We capture this intuition with *connectivity of a similarity graph*. We define the *similarity graph* of a set  $\mathbf{R}$  of records as an undirected graph, where each node represents a record in  $\mathbf{R}$ , and an edge indicates high similarity between the connected records (we describe later what we mean by *high similarity*). Figure 1 shows the similarity graph for the motivating example.

Each core would correspond to a connected sub-graph of the similarity graph. We wish such a sub-graph to be *robust* such that




Figure 1: Similarity graph for records in Table 2.

even if we remove a few nodes the sub-graph is still connected; in other words, even if there are some erroneous records, without them we still have enough evidence showing that the rest of the records should belong to the same group. The formal definition goes as follows.

**DEFINITION 4.1** ( $k$ -ROBUSTNESS). *A graph  $G$  is  $k$ -robust if after removing arbitrary  $k$  nodes and edges to these nodes,  $G$  is still connected. A clique or a single node is  $k$ -robust for any  $k$ . □*

In Figure 1, the subgraph with nodes  $r_1 - r_7$  is 2-robust. That with  $r_{11} - r_{18}$  is not 1-robust, as removing  $r_{13}$  can disconnect it.

According to the definition, we can partition the similarity graph into a set of  $k$ -robust subgraphs. As we do not wish to split any core unnecessarily, we require the *maximal  $k$ -robust partitioning*:

**DEFINITION 4.2** (MAXIMAL  $k$ -ROBUST PARTITIONING). *Let  $G$  be a similarity graph. A partitioning of  $G$  is a maximal  $k$ -robust partitioning if it satisfies the following properties.*

1. *Each node belongs to one and only one partition.*
2. *Each partition is  $k$ -robust.*
3. *The result of merging any partitions is not  $k$ -robust.* □

Note that a data set can have more than one maximal  $k$ -robust partitioning. Consider  $r_{11} - r_{18}$  in Figure 1. There are three maximal 1-robust partitionings:  $\{\{r_{11}\}, \{r_{12}, r_{14} - r_{15}\}, \{r_{13}, r_{16} - r_{18}\}\}$ ;  $\{\{r_{11} - r_{12}\}, \{r_{14} - r_{15}\}, \{r_{13}, r_{16} - r_{18}\}\}$ ; and  $\{\{r_{11} - r_{15}\}, \{r_{16} - r_{18}\}\}$ . If we treat each partitioning as a possible world, records that belong to the same partition in all possible worlds have high probability to belong to the same group and so form a core. Accordingly, we define a core as follows and can prove its  $k$ -robustness.

**DEFINITION 4.3** ( $k$ -CORE). *Let  $\mathbf{R}$  be a set of records and  $G$  be the similarity graph of  $\mathbf{R}$ . The records that belong to the same subgraph in every maximal  $k$ -robust partitioning of  $G$  form a  $k$ -core of  $\mathbf{R}$ . A core contains at least 2 records.* □

**PROPERTY 4.4.** *A  $k$ -core is  $k$ -robust.* □

**PROOF.** If a  $k$ -core  $C_r$  of  $G$  is not  $k$ -robust, there exists a maximal  $k$ -robust partitioning in  $G$ , where two nodes  $r$  and  $r'$  in  $C_r$  are in different partitions of this partitioning (proved by Lemma 4.19). This conflicts with the fact that records in  $C_r$  belong to the same partition in every maximal  $k$ -robust partitioning of  $G$ . Therefore, a  $k$ -core is  $k$ -robust. □

**EXAMPLE 4.5.** *Consider Figure 1 and assume  $k = 1$ . There are two connected sub-graphs. For records  $r_1 - r_7$ , the subgraph is 1-robust, so they form a 1-core. For records  $r_{11} - r_{18}$ , there are three maximal 1-robust partitionings for the subgraph, as we have shown. Two subsets of records belong to the same subgraph in each partitioning:  $\{r_{14} - r_{15}\}$  and  $\{r_{16} - r_{18}\}$ ; they form 2 1-cores. □*

**Table 3: Simplified inverted index for the similarity graph in Figure 1.**

Record	V-Cliques	Represent
$r_{1/2}$	$C_1$	$r_1 - r_2$
$r_3$	$C_1, C_2$	$r_3$
$r_4$	$C_1, C_2$	$r_4$
$r_5$	$C_1, C_2$	$r_5$
$r_{6/7}$	$C_2$	$r_6 - r_7$
$r_{11}$	$C_3$	$r_{11}$
$r_{12}$	$C_3, C_4$	$r_{12}$
$r_{13}$	$C_3, C_4, C_5$	$r_{13}$
$r_{14/15}$	$C_4$	$r_{14} - r_{15}$
$r_{16/17/18}$	$C_5$	$r_{16} - r_{18}$

## 4.2 Constructing similarity graphs

Generating the cores requires analysis on the similarity graph. Even after blocking, a block can contain tens of thousands of records, so it is not scalable to compare every pair of records in the same block and create edges accordingly. We next describe how we construct and represent the similarity graph in a scalable way.

We add an edge between two records if they have the same value for each common-value attribute and share at least one value on a dominant-value attribute<sup>2</sup>; our experiments show advantages of this method over other edge-adding strategies (Section 6.2.1). All records that share values on the common-value attributes and share the same value on a dominant-value attribute form a clique, which we call a *v-clique*. We can thus represent the graph with a set of v-cliques, denoted by  $\mathbf{C}$ ; for example, the graph in Figure 1 can be represented by 5 v-cliques ( $C_1 - C_5$ ). In addition, we maintain an *inverted index*  $\bar{L}$ , where each entry corresponds to a record  $r$  and contains the v-cliques that  $r$  belongs to. Whereas the size of the similarity graph can be quadratic in the number of the nodes, the size of the inverted index is only linear in that number. The inverted index also makes it easy to find *adjacent v-cliques* (i.e., v-cliques that share nodes), as they appear in the same entry.

Graph construction is then reduced to v-clique finding, which can be done by scanning values of dominant-value attributes. In this process, we wish to prune a v-clique if it is a sub-clique of another one. Pruning by checking every pair of v-cliques can be very expensive since the number of v-cliques is also huge. Instead, we do it together with v-clique finding. Specifically, our algorithm GRAPHCONSTRUCTION takes  $\mathbf{R}$  as input and outputs  $\mathbf{C}$  and  $\bar{L}$ . We start with  $\mathbf{C} = \bar{L} = \emptyset$ . For each value  $v$  of a dominant-value attribute, we denote the set of records with  $v$  by  $\bar{R}_v$  and do the following.

1. Initialize the v-cliques for  $v$  as  $\mathbf{C}_v = \emptyset$ . Add a single-record cluster for each record  $r \in \bar{R}_v$  to a working set  $\bar{T}$ . Mark each cluster as “unchanged”.
2. For each  $r \in \bar{R}_v$ , scan  $\bar{L}$  and consider each v-clique  $C \in \bar{L}(r)$  that has not been considered yet. For all records in  $C \cap \bar{R}_v$ , merge their clusters. Mark the merged cluster as “changed” if the result is not a proper sub-clique of  $\bar{C}$ . If  $C \subseteq \bar{R}_v$ , remove  $C$  from  $\mathbf{C}$ . This step removes the v-cliques that must be sub-cliques of those we will form next.
3. For each cluster  $C \in \bar{T}$ , if there exists  $C' \in \mathbf{C}_v$  such that  $C$  and  $C'$  share the same value for each common-value attribute, remove  $C$  and  $C'$  from  $\bar{T}$  and  $\mathbf{C}_v$  respectively, add  $C \cup C'$  to  $\bar{T}$  and mark it as “changed”; otherwise, move  $C$  to  $\mathbf{C}_v$ . This step merges clusters that share values on common-value attributes. At the end,  $\mathbf{C}_v$  contains the v-cliques with value  $v$ .

<sup>2</sup>In practice, we require only highly similar values for common-value attributes and apply the transitive rule on similarity (i.e., if  $v_1$  and  $v_2$  are highly similar, and so are  $v_2$  and  $v_3$ , we consider  $v_1$  and  $v_3$  highly similar).

4. Add each v-clique with mark “changed” in  $\mathbf{C}_v$  to  $\mathbf{C}$  and update  $\bar{L}$  accordingly. The marking prunes size-1 v-cliques and the sub-cliques of those already in  $\mathbf{C}$ .

**PROPOSITION 4.6.** *Let  $\mathbf{R}$  be a set of records. Denote by  $n(r)$  the number of values on dominant-value attributes from  $r \in \mathbf{R}$ . Let  $n = \sum_{r \in \mathbf{R}} n(r)$  and  $m = \max_{r \in \mathbf{R}} n(r)$ . Let  $s$  be the maximum v-clique size. Algorithm GRAPHCONSTRUCTION (1) runs in time  $O(ns(m + s))$ , (2) requires space  $O(n)$ , and (3) its result is independent of the order in which we consider the records.*  $\square$

**PROOF.** We first prove that GRAPHCONSTRUCTION runs in time  $O(ns(m + s))$ . Step 2 of the algorithm takes in time  $O(nsm)$ , where it takes in time  $O(ns)$  to scan all records for a dominant-value attribute, and a record can be scanned maximally  $m$  times. Step 3 takes in time  $O(ns^2)$ . Thus, the algorithm runs in time  $O(ns(m + s))$ .

We next prove that GRAPHCONSTRUCTION requires space  $O(n)$ . For each value  $v$  of a dominate-value attribute, the algorithm keeps three data sets:  $\bar{L}$  that takes in space  $O(n)$ ,  $\mathbf{C}_v$  and  $\bar{T}$  that require space in total no greater than  $O(|\mathbf{R}|)$ . Since  $O(n) \geq O(|\mathbf{R}|)$ , the algorithm requires space  $O(n)$ .

We now prove that the result of GRAPHCONSTRUCTION is order independent. Given  $\bar{L}$  and  $\bar{R}_v$ , Step 2 scan  $\bar{L}$  and apply transitive rule to merge clusters of records in  $C \cap \bar{R}_v$ , for each v-clique  $C \in \bar{L}$ . The process is independent from the order in which we consider the records in  $\bar{R}_v$ . The order independence of the result in Step 3 is proven in [2]. Therefore, the final result is independent from the order in which we consider the records.  $\square$

**EXAMPLE 4.7.** *Consider graph construction for records in Table 2. Figure 1 shows the similarity graph and Table 3(a) shows the inverted list. We focus on records  $r_1 - r_8$  for illustration.*

*First,  $r_1 - r_5$  share the same name and phone number 808, so we add v-clique  $C_1 = \{r_1 - r_5\}$  to  $\mathbf{C}$ . Now consider URL homedepot where  $\bar{R}_v = \{r_3 - r_8\}$ . Step 1 generates 6 clusters, each marked “unchanged”, and  $\bar{T} = \{\{r_3\}, \dots, \{r_8\}\}$ . Step 2 looks up  $\bar{L}$  for each record in  $\bar{R}_v$ . Among them,  $r_3 - r_5$  belong to v-clique  $C_1$ , so it merges their clusters and marks the result  $\{r_3 - r_5\}$  “unchanged” ( $\{r_3 - r_5\} \subset C_1$ ); then,  $\bar{T} = \{\{r_3 - r_5\}, \{r_6\}, \{r_7\}, \{r_8\}\}$ . Step 3 compares these clusters and merges the first three as they share the same name, marking the result as “changed”. At the end,  $\mathbf{C}_v = \{\{r_3 - r_7\}, \{r_8\}\}$ . Finally, Step 4 adds  $\{r_3 - r_7\}$  to  $\mathbf{C}$  and discards  $\{r_8\}$  since it is marked “unchanged”.  $\square$*

Given the sheer number of records in  $\mathbf{R}$ , the inverted index can still be huge. In fact, according to the following theorem, records in the same v-clique but not any other v-clique must belong to the same core, so we do not need to distinguish them. Thus, we simplify the inverted index such that for each v-clique we keep only a *representative* for nodes belonging only to this v-clique. Table 3 shows the simplified index for the similarity graph in Figure 1.

**THEOREM 4.8.** *Let  $G$  be a similarity graph and  $G'$  be a graph derived from  $G$  by merging nodes that belong to only one and the same v-clique. Two nodes belong to the same core of  $G'$  if and only if they belong to the same core of  $G$ .*  $\square$

**PROOF.** We need to prove that (1) if two nodes  $r$  and  $r'$  belong to the same core in  $G'$ , they are in the same core of  $G$ , and (2) if two nodes  $r$  and  $r'$  belong to the same core of  $G$ , they are in the same core of  $G'$ .

We first prove that if two nodes  $r$  and  $r'$  belong to the same core in  $G'$ , they are in the same core of  $G$ . Suppose there does not exist any core in  $G$  that contains both  $r$  and  $r'$ . It means that

there exists a maximal  $k$ -robust partitioning in  $G$ , where  $r$  and  $r'$  are in different partitions. Let  $P$  be such a partitioning of  $G$  and we consider partitioning  $P'$  of  $G'$ , where each pair of nodes in the same partition  $C$  of  $P$  are in the same partition  $C'$  of  $P'$  and vice versa. We prove that  $P'$  is a maximal  $k$ -robust partitioning in  $G'$ . (1) It is obvious that each node in  $P'$  belongs to one and only one partition. (2) For each partition  $C'$  in  $P'$ , removing any  $k$  nodes in  $C'$  is equivalent to removing  $n + m$  nodes in  $C$ , where  $n$  nodes belong to more than one  $v$ -cliques in  $C$ ,  $m$  nodes belong to single  $v$ -cliques in  $C$ , and  $n \leq k$ . Since removing  $m$  nodes that belong to single  $v$ -cliques do not disconnect  $C$  and we know  $n \leq k$ , removing the  $n + m$  nodes does not disconnect  $C$ . It in turn proves that removing  $k$  nodes in  $C'$  does not disconnect  $C'$ , and  $C'$  is  $k$ -robust. (3) Similarly, we have that the result of merging any partitions in  $P'$  is not  $k$ -robust. Therefore,  $P'$  is a maximal  $k$ -robust partitioning in  $G'$ . Given that  $r$  and  $r'$  are in different partitions of  $P'$ , there does not exist a core of  $G'$  that contains both  $r$  and  $r'$ . This conflicts with the fact that  $r$  and  $r'$  belong to the same core in  $G'$ , and further proves that  $r$  and  $r'$  are in the same core of  $G$ .

We next prove that if two nodes  $r$  and  $r'$  belong to the same core of  $G$ , they are in the same core of  $G'$ . Suppose there does not exist any core in  $G'$  that contains both  $r$  and  $r'$ . It means that there exists a maximal  $k$ -robust partitioning in  $G'$ , where  $r$  and  $r'$  are in different partitions. Let  $P'$  be such a partitioning of  $G'$  and we consider partitioning  $P$  of  $G$ , where each pair of nodes in the same partition  $C'$  of  $P'$  are in the same partition  $C$  of  $P$  and vice versa. In similar ways as above, we have that  $P$  is a maximal  $k$ -robust partitioning in  $G$ . Given that  $r$  and  $r'$  are in different partitions of  $P$ , there does not exist a core of  $G$  that contains both  $r$  and  $r'$ . This conflicts with the fact that  $r$  and  $r'$  belong to the same core in  $G$ , and further proves that  $r$  and  $r'$  are in the same core of  $G'$ .  $\square$

**Case study:** On a data set with 18M records (described in Section 6), our graph-construction algorithm finished in 1.9 hours. The original similarity graph contains 18M nodes and 4.2B edges. The inverted index is of size 89MB, containing 3.8M entries, each associated with at most 8  $v$ -cliques; in total there are 1.2M  $v$ -cliques. The simplified inverted index is of size 34MB, containing 1.5M entries, where an entry can represent up to 11K records. Therefore, the simplified inverted index reduces the size of the similarity graph by 3 orders of magnitude.

### 4.3 Identifying cores

We solve the core-identification problem by reducing it to a Max-flow/Min-cut Problem. However, computing the max flow for a given graph  $G$  and a source-destination pair takes time  $O(|G|^{2.5})$ , where  $|G|$  denotes the number of nodes in  $G$ ; even the simplified inverted index can still contain millions of entries, so it can be very expensive. We thus first merge certain  $v$ -cliques according to a sufficient (but not necessary) condition for  $k$ -robustness and consider them as a whole in core identification; we then split the graph into subgraphs according to a necessary (but not sufficient) condition for  $k$ -robustness. We apply reduction only on the resulting subgraphs, which are substantially smaller as we show at the end of this section. Section 4.3.1 describes screening before reduction, Section 4.3.2 describes the reduction, and Section 4.3.3 gives the full algorithm, which iteratively applies screening and the reduction.

#### 4.3.1 Screening

A graph can be considered as a union of  $v$ -cliques, so essentially we need to decide if a union of  $v$ -cliques is  $k$ -robust. First, we can prove the following sufficient condition for  $k$ -robustness.




Figure 2: Two example graphs.

**THEOREM 4.9 (( $K + 1$ )-CONNECTED CONDITION).** *Let  $G$  be a graph consisting of a union  $Q$  of  $v$ -cliques. If for every pair of  $v$ -cliques  $C, C' \in Q$ , there is a path of  $v$ -cliques between  $C$  and  $C'$  and every pair of adjacent  $v$ -cliques on the path share at least  $k + 1$  nodes, graph  $G$  is  $k$ -robust.*  $\square$

**PROOF.** Given Menger's Theorem [3], graph  $G$  is  $k$ -robust if for any pair of nodes  $r, r'$  in  $G$ , there exists at least  $k + 1$  independent paths that do not share any nodes other than  $r, r'$  in  $G$ . We now prove that for any pair of nodes  $r, r'$  in graph  $G$  that satisfies  $(k + 1)$ -connected condition, there exists at least  $k + 1$  independent paths between  $r, r'$ . We consider two cases, 1)  $r, r'$  are adjacent such that there exists a  $v$ -clique in  $G$  that contains  $r, r'$ ; 2)  $r, r'$  are not adjacent such that there exists no  $v$ -clique in  $G$  that contains  $r, r'$ .

We first consider Case 1 where there exists a  $v$ -clique  $C$  containing  $r, r'$ . Since each  $v$ -clique in  $G$  has more than  $k + 1$  nodes, there exist at least  $k$  2-length paths and one 1-length path between  $r, r' \in C$ . It proves that there exists at least  $k + 1$  independent paths between  $r$  and  $r'$ .

We next consider Case 2 where there exists no  $v$ -clique containing  $r, r'$  in  $G$ . Suppose  $r \in C, r' \in C'$ , where  $C, C'$  are different  $v$ -cliques in  $G$ . Since there exists a path of  $v$ -cliques between  $C$  and  $C'$  where every pair of adjacent  $v$ -cliques in the path share at least  $k + 1$  nodes, there exists at least  $k + 1$  independent paths between  $r$  and  $r'$ .

Given the above two cases, we have that there exist at least  $k + 1$  independent paths between every pair of nodes in  $G$ , therefore  $G$  is  $k$ -robust.  $\square$

We call a single  $v$ -clique or a union of  $v$ -cliques that satisfy the  $(k + 1)$ -connected condition a  $(k + 1)$ -connected  $v$ -union. A  $(k + 1)$ -connected  $v$ -union must be  $k$ -robust but not vice versa. In Figure 1, subgraph  $\{r_{11} - r_{18}\}$  is a 3-connected  $v$ -union, because the only two  $v$ -cliques,  $C_1$  and  $C_2$ , share 3 nodes. Indeed, it is 2-robust. On the other hand, graph  $G_1$  in Figure 2 is 2-robust but not 3-connected (there are 4  $v$ -cliques, where each pair of adjacent  $v$ -cliques share only 1 or 2 nodes). Accordingly, we can consider a  $v$ -union as a whole in core identification.

Next, we present a necessary condition for  $k$ -robustness.

**THEOREM 4.10 (( $K + 1$ )-OVERLAP CONDITION).** *Graph  $G$  is  $k$ -robust only if for every  $(k + 1)$ -connected  $v$ -union  $Q \in G$ ,  $Q$  shares at least  $k + 1$  common nodes with the subgraph consisting of the rest of the  $v$ -unions.*  $\square$

**PROOF.** We prove that if graph  $G$  contains a  $(k + 1)$ -connected  $v$ -union  $Q$  that shares at most  $k$  common nodes with the rest of the graph,  $G$  is not  $k$ -robust. Since  $Q$  shares at most  $k$  common nodes with the subgraph consisting of the rest of the  $v$ -unions, removing the common nodes will disconnect  $Q$  from  $G$ , it proves that  $G$  is not  $k$ -robust. Thus,  $(k + 1)$ -overlap condition holds.  $\square$

We call a graph  $G$  that satisfies the  $(k + 1)$ -overlap condition a  $(k + 1)$ -overlap graph. A  $k$ -robust graph must be a  $(k + 1)$ -overlap graph but not vice versa. In Figure 1, subgraph  $\{r_{11} - r_{18}\}$  is not a 2-overlap graph, because there are two 2-connected  $v$ -unions,  $\{r_{11} - r_{15}\}$  and  $\{r_{13}, r_{16} - r_{18}\}$ , but they share only one node;



indeed, the subgraph is not 1-robust. On the other hand, graph  $G_2$  in Figure 2 satisfies the 3-overlap condition, as it contains four 3-connected v-unions (actually four v-cliques),  $Q_1 - Q_4$ , and each v-union shares 3 nodes in total with the others; however, it is not 2-robust (removing  $r_3$  and  $r_4$  disconnects it). Accordingly, for  $(k + 1)$ -overlap graphs we still need to check  $k$ -robustness by reduction to a Max-flow Problem.

Now the problem is to find  $(k + 1)$ -overlap subgraphs. Let  $G$  be a graph where a  $(k + 1)$ -connected v-union overlaps with the rest of the v-unions on no more than  $k$  nodes. We split  $G$  by removing these overlapping nodes. For subgraph  $\{r_{11} - r_{18}\}$  in Figure 1, we remove  $r_{13}$  and obtain two subgraphs  $\{r_{11} - r_{12}, r_{14} - r_{15}\}$  and  $\{r_{16} - r_{18}\}$  (recall from Example 4.5 that  $r_{13}$  cannot belong to any core). Note that the result subgraphs may not be  $(k + 1)$ -overlap graphs (e.g.,  $\{r_{11} - r_{12}, r_{14} - r_{15}\}$  contains two v-unions that share only one node), so we need to further screen them.

We now describe our screening algorithm, SCREEN (details in Algorithm 1), which takes a graph  $G$ , represented by  $\mathbf{C}$  and  $\bar{L}$ , as input, finds  $(k + 1)$ -connected v-unions in  $G$  and meanwhile decides if  $G$  is a  $(k + 1)$ -overlap graph. If not, it splits  $G$  into subgraphs for further examination.

1. If  $G$  contains a single node, output it as a core if the node represents multiple records that belong only to one v-clique.
  2. For each v-clique  $C \in \mathbf{C}$ , initialize a v-union. We denote the set of v-unions by  $\bar{Q}$ , the v-union that  $C$  belongs to by  $Q(C)$ , and the overlapping nodes of  $C$  and  $C'$  by  $\bar{B}(C, C')$ .
  3. For each v-clique  $C \in \mathbf{C}$ , we merge v-unions as follows.
    - (a) For each record  $r \in C$  that has not been considered, for every pair of v-cliques  $C_1$  and  $C_2$  in  $r$ 's index entry, if they belong to different v-unions, add  $r$  to overlap  $\bar{B}(C_1, C_2)$ .
    - (b) For each v-union  $Q \neq Q(C)$  where there exist  $C_1 \in Q$  and  $C_2 \in Q(C)$  such that  $|\bar{B}(C_1, C_2)| \geq k + 1$ , merge  $Q$  and  $Q(C)$ .
- At the end,  $\bar{Q}$  contains all  $(k + 1)$ -connected v-unions.
4. For each v-union  $Q \in \bar{Q}$ , find its border nodes as  $\bar{B}(Q) = \cup_{C \in Q, C' \notin Q} \bar{B}(C, C')$ . If  $|\bar{B}(Q)| \leq k$ , split the subgraph it belongs to, denoted by  $G(Q)$ , into two subgraphs  $Q \setminus \bar{B}(Q)$  and  $G(Q) \setminus Q$ .
  5. Return the remaining subgraphs.

**PROPOSITION 4.11.** *Denote by  $|\bar{L}|$  the number of entries in input  $\bar{L}$ . Let  $m$  be the maximum number of values from dominant-value attributes of a record, and  $a$  be the maximum number of adjacent v-unions that a v-union has. Algorithm SCREEN finds  $(k + 1)$ -overlap subgraphs in time  $O((m^2 + a) \cdot |\bar{L}|)$  and the result is independent of the order in which we examine the v-cliques.  $\square$*

**PROOF.** We first prove the time complexity of SCREEN. It takes in time  $O(m^2 |\bar{L}|)$  to scan all entries in  $\bar{L}$  and find common nodes between each pair of adjacent v-cliques (Step 3(a)). It takes in time  $O(a |\mathbf{C}|)$  to merge v-unions, where  $|\mathbf{C}|$  is the number of v-cliques in  $G$  (Step 3(b)). Since  $|\mathbf{C}| < |\bar{L}|$ , the algorithm runs in time  $O(m^2 + a) \cdot |\bar{L}|$ .

We next prove that the result of *Screen* is independent of the order in which we examine the v-cliques, that is, 1) finding all maximal  $(k + 1)$ -connected v-unions in  $G$  is order independent; 2) removing all nodes in  $\bar{B}(Q)$  from  $G$  where  $|\bar{B}(Q)| \leq k$  is order independent.

Consider order independency of finding all v-unions in  $G$ . To find all v-unions in  $G$  is conceptually equivalent to find all connected components in an abstract graph  $G_A$ , where each node in

---

**Algorithm 1** SCREENING( $G, \bar{C}, \bar{L}, k$ )

---

**Input:**  $G$ : Simplified similarity graph.

$\bar{C}$ : Set of  $k$ -cores.

$\bar{L}$ : Inverted list of the similarity graph.

$k$ : Robustness requirement.

**Output:**  $\bar{G}$  Set of subgraphs in  $G$ .

```

1: if  $G$  contains a single node  $r$  then
2:   if  $r$  represent multiple records then
3:     add  $r$  to  $\bar{C}$ .
4:   end if
5:   return  $\bar{G} = \phi$ .
6: else
7:   initialize v-union  $Q(C)$  for each v-clique  $C$  and add  $Q(C)$  to  $\bar{Q}$ .
8:   // find v-union
9:   for each v-clique  $C \in G$  do
10:    for each record  $r \in C$  that is not proceeded do
11:      for each v-clique pair  $C_1, C_2 \in \bar{L}(r)$  do
12:        if  $C_1, C_2$  are in different v-unions then
13:          add  $r$  to overlap  $\bar{B}(Q(C_1), Q(C_2))$ .
14:        end if
15:      end for
16:    end for
17:    for each v-union  $Q$  where  $|\bar{B}(Q, Q(C))| \geq k$  do
18:      merge  $Q$  and  $Q(C)$  as  $Q_m$ .
19:      for each v-union  $Q' \neq Q, Q' \neq Q(C)$  do
20:        set  $\bar{B}(Q', Q_m) = \bar{B}(Q', Q) \cup \bar{B}(Q', Q(C))$ 
21:      end for
22:    end for
23:  end for
24:  // screening
25:  for each v-union  $Q \in \bar{Q}$  do
26:    compute  $\bar{B}(Q) = \cup_{Q' \in \bar{Q}} \bar{B}(Q, Q')$ .
27:    if  $|\bar{B}(Q)| < k$  then
28:      add subgraphs  $Q \setminus \bar{B}(Q)$  and  $G(Q) \setminus Q$  into  $\bar{G}$ 
29:    end if
30:  end for
31: end if
32: return  $\bar{G}$ ;

```

---

$G_A$  is a v-clique in  $G$  and two nodes in  $G_A$  are connected if the two corresponding v-cliques share more than  $k$  nodes. SCREEN checks whether each node in  $G$  is a common node between two v-cliques (Step 3(a)), and if two cliques share more than  $k$  nodes, merges their v-unions (Step 3(b)), which is equivalent to connect two nodes in  $G_A$ . Once all nodes in  $G$  is scanned, all edges in  $G_A$  are added, and the order in which we examine nodes in  $G$  is independent from the structure of  $G_A$  and the connected components in  $G_A$ . Therefore, finding all v-unions in  $G$  is order independent.

Consider order independency of removing nodes in  $G$ . Suppose  $Q_1, Q_2, \dots, Q_m, m > 0$  are all v-unions in  $G$  with  $|\bar{B}(Q_i)| \leq k, i \in [1, m]$ . Since  $G$  is finite,  $Q_i$  is finite and unique; thus, removing all nodes in  $\bar{B}(Q)$  from  $G$  where  $|\bar{B}(Q)| \leq k$  is order independent.  $\square$

Note that  $m$  and  $a$  are typically very small, so SCREEN is basically linear in the size of the inverted index. Finally, we have results similar to Theorem 4.8 for v-unions, so we can further simplify the graph by keeping for each v-union a single representative for all nodes that only belong to it. Each result  $k$ -overlap subgraph is typically very small.

EXAMPLE 4.12. Consider Table 3 as input and  $k = 1$ . Step 2 creates five v-unions  $Q_1 - Q_5$  for the five v-cliques in the input.

Step 3 starts with v-clique  $C_1$ . It has 4 nodes (in the simplified inverted index), among which 3 are shared with  $C_2$ . Thus,  $\bar{B}(C_1, C_2) = \{r_3 - r_5\}$  and  $|\bar{B}(C_1, C_2)| \geq 2$ , so we merge  $Q_1$  and  $Q_2$  into  $Q_{1/2}$ . Examining  $C_2$  reveals no other shared node.

Step 3 then considers v-clique  $C_3$ . It has three nodes, among which  $r_{12} - r_{13}$  are shared with  $C_4$  and  $r_{13}$  is also shared with  $C_5$ . Thus,  $\bar{B}(C_3, C_4) = \{r_{12} - r_{13}\}$  and  $\bar{B}(C_3, C_5) = \{r_{13}\}$ . We merge  $Q_3$  and  $Q_4$  into  $Q_{3/4}$ . Examining  $C_4$  and  $C_5$  reveals no other shared node. We thus obtain three 2-connected v-unions:  $\bar{Q} = \{Q_{1/2}, Q_{3/4}, Q_5\}$ .

Step 4 then considers each v-union. For  $Q_{1/2}$ ,  $\bar{B}(Q_{1/2}) = \emptyset$  and we thus split subgraph  $Q_{1/2}$  out and merge all of its nodes to one  $r_{1/\dots/7}$ . For  $Q_{3/4}$ ,  $\bar{B}(Q_{3/4}) = \{r_{13}\}$  so  $|\bar{B}(Q_{3/4})| < 2$ . We split  $Q_{3/4}$  out and obtain  $\{r_{11} - r_{12}, r_{14/15}\}$  ( $r_{13}$  is excluded). Similar for  $Q_5$  and we obtain  $\{r_{16/17/18}\}$ . Therefore, we return three subgraphs for further screening.  $\square$

### 4.3.2 Reduction

Intuitively, a graph  $G(V, E)$  is  $k$ -robust if and only if between any two nodes  $a, b \in V$ , there are more than  $k$  paths that do not share any node except  $a$  and  $b$ . We denote the number of non-overlapping paths between nodes  $a$  and  $b$  by  $\kappa(a, b)$ . We can reduce the problem of computing  $\kappa(a, b)$  into a Max-flow Problem.

For each input  $G(V, E)$  and nodes  $a, b$ , we construct the (directed) flow network  $G'(V', E')$  as follows.

1. Node  $a$  is the source and  $b$  is the sink (there is no particular order between  $a$  and  $b$ ).
2. For each  $v \in V, v \neq a, v \neq b$ , add two nodes  $v', v''$  to  $V'$ , and two directed edges  $(v', v''), (v'', v')$  to  $E'$ . If  $v'$  represents  $n$  nodes, the edge  $(v', v'')$  has weight  $n$ , and the edge  $(v'', v')$  has weight  $\infty$ .
3. For each edge  $(a, v) \in E$ , add edge  $(a, v')$  to  $E'$ ; for each edge  $(u, b) \in E$ , add edge  $(u'', b)$  to  $E'$ ; for each other edge  $(u, v) \in E$ , add two edges  $(u'', v')$  and  $(v'', u')$  to  $E'$ . Each edge has capacity  $\infty$ .

LEMMA 4.13. The max flow from source  $a$  to sink  $b$  in  $G'(V', E')$  is equivalent to  $\kappa(a, b)$  in  $G(V, E)$ .  $\square$

PROOF. According to Menger's Theorem [3], the minimum number of nodes whose removal disconnects  $a$  and  $b$ , that is  $\kappa(a, b)$ , is equal to the maximum number of independent paths between  $a$  and  $b$ . The authors in [9] proves that the maximum number of independent paths between  $a$  and  $b$  in an undirected graph  $G(V, E)$  is equivalent to the maximal value of flow from  $a$  to  $b$  or the minimal capacity of an  $a - b$  cut, the set of nodes such that any path from  $a$  to  $b$  contains a member of the cut, in  $G'(V', E')$ .  $\square$

EXAMPLE 4.14. Consider nodes  $r_1$  and  $r_6$  of graph  $G_2$  in Figure 2. Figure 3 shows the corresponding flow network, where the dash line (across edges  $(r_3', r_3'')$ ,  $(r_4', r_4'')$ ) in the figure cuts the flow from  $r_1$  to  $r_6$  with a minimum cost of 2. The max flow/min cut has value 2. Indeed,  $\kappa(r_1, r_6) = 2$ .  $\square$

Recall that in a  $(k + 1)$ -connected v-union, between each pair of nodes there are at least  $k + 1$  paths. Thus, if (1)  $\kappa(a, b) = k + 1$ , (2)  $a$  and  $b$  belong to different v-unions, and (3)  $a$  and  $a'$  belong to the same v-union, we must have  $\kappa(a', b) \geq k + 1$ . We thus have the following sufficient and necessary condition for  $k$ -robustness.

THEOREM 4.15 (MAX-FLOW CONDITION). Let  $G(V, E)$  be an input similarity graph. Graph  $G$  is  $k$ -robust if and only if for




Figure 3: Flow network for  $G_2$  in Figure 2.

every pair of adjacent  $(k + 1)$ -connected v-unions  $Q$  and  $Q'$ , there exist two nodes  $a \in Q \setminus Q'$  and  $b \in Q' \setminus Q$  such that the max flow from  $a$  to  $b$  in the corresponding flow network is at least  $k + 1$ .  $\square$

PROOF. According to Menger's Theorem [3],  $\kappa(a, b)$  in  $G$  is equivalent to the max-flow from  $a$  to  $b$  in the corresponding flow network. We need to prove that graph  $G$  is  $k$ -robust if and only if for each pair of adjacent  $(k + 1)$ -connected v-unions  $Q$  and  $Q'$ , there exists two nodes  $a \in Q \setminus Q'$  and  $b \in Q' \setminus Q$  such that  $\kappa(a, b) \geq k + 1$ .

We first prove that if  $G$  is  $k$ -robust, for each pair of adjacent  $(k + 1)$ -connected v-unions  $Q$  and  $Q'$ , there exists two nodes  $a \in Q \setminus Q'$  and  $b \in Q' \setminus Q$  such that  $\kappa(a, b) \geq k + 1$ . Since  $G$  is  $k$ -robust, for each pair of nodes  $a$  and  $b$  in  $G$ , we have  $\kappa(a, b) \geq k + 1$ .

We next prove that if  $G$  is not  $k$ -robust, there exists a pair of adjacent  $(k + 1)$ -connected v-unions  $Q$  and  $Q'$  such that for each pair of nodes  $a \in Q \setminus Q'$  and  $b \in Q' \setminus Q$ , we have  $\kappa(a, b) < k + 1$ . Since  $G$  is not  $k$ -robust, there exists a separator  $\bar{S}$ , a set of nodes in  $G$  with size no greater than  $k$  whose removal disconnects  $G$  into two sub-graphs  $\bar{X}$  and  $\bar{Y}$ . Suppose  $Q$  and  $Q'$  are two v-unions in  $G$  such that  $Q \subseteq \bar{X}, Q' \subseteq \bar{Y}$  and  $Q \cap Q' \neq \emptyset$ . For each pair of nodes  $a \in Q \setminus Q'$  and  $b \in Q' \setminus Q$ , we have  $a \in \bar{X}$  and  $b \in \bar{Y}$ , and removing the set of nodes in  $\bar{S}$  disconnects  $a$  and  $b$ ; thus  $\kappa(a, b) < k + 1$ .

The above two cases proves that graph  $G$  is  $k$ -robust if and only if for every pair of adjacent  $(k + 1)$ -connected v-unions  $Q$  and  $Q'$ , there exist two nodes  $a \in Q \setminus Q'$  and  $b \in Q' \setminus Q$  such that  $\kappa(a, b) \geq k + 1$ , i.e. the max flow from  $a$  to  $b$  in the corresponding flow network is at least  $k + 1$ .  $\square$

If a graph  $G$  is not  $k$ -robust, we shall split it into subgraphs for further processing. In the corresponding flow network, each edge in the minimum cut must be between a pair of nodes derived from the same node in  $G$  (other edges have capacity  $\infty$ ). These nodes cannot belong to any core and we use them as separator nodes, denoted by  $\bar{S}$ . Suppose the separator separates  $G$  into  $\bar{X}$  and  $\bar{Y}$  (there can be more subgraphs); we return  $\bar{X} \cup \bar{S}$  and  $\bar{Y} \cup \bar{S}$ .

Note that we need to include  $\bar{S}$  in both sub-graphs to maintain the integrity of each v-union. To understand why, consider  $G_2$  in Figure 2 where  $\bar{S} = \{r_3, r_4\}$ . According to the definition, there is no 2-core. If we split  $G_2$  into  $\{r_1 - r_2\}$  and  $\{r_5 - r_6\}$  (without including  $\bar{S}$ ), both subgraphs are 2-robust and we would return them as 2-cores. The problem happens because v-cliques  $Q_1 - Q_4$  "disappear" after we remove the separators  $r_3$  and  $r_4$ . Thus, we should split  $G_2$  into  $\{r_1 - r_4\}$  and  $\{r_3 - r_6\}$  instead and that would further trigger splitting on both subgraphs. Eventually we wish to exclude the separator nodes from any core, so we mark them as "separators" and exclude them from the returned cores.

Algorithm SPLIT (details in Algorithm 2) takes a  $(k + 1)$ -overlap subgraph  $G$  as input and decides if  $G$  is  $k$ -robust. If not, it splits  $G$  into subgraphs on which we will then re-apply screening.

1. For each pair of adjacent  $(k + 1)$ -connected v-unions  $Q, Q' \in G$ , find  $a \in Q \setminus Q', b \in Q' \setminus Q$ . Construct flow network



---

**Algorithm 2** SPLIT( $G, \bar{C}, k$ )

---

**Input:**  $G$ : Simplified similarity graph.

$\bar{C}$ : Set of cores.

$k$ : Robustness requirement.

**Output:**  $\bar{G}$  Set of subgraphs in  $G$ .

```
1: for each adjacent  $(k + 1)$ -connected v-unions  $Q, Q'$  do
2:   find a pair of nodes  $a \in Q \setminus Q', b \in Q' \setminus Q$ .
3:   construct flow-network  $G'$  and compute  $\kappa(a, b)$  by Ford &
   Fulkerson Algorithm.
4:   if  $\kappa(a, b) \leq k$  then
5:     get separator  $\bar{S}$  from  $G'$  and remove  $\bar{S}$  from  $G$  to obtain
     disconnected subgraphs; mark  $\bar{S}$  as “separator” and add it
     to each subgraph in  $G$ .
6:   return the set  $\bar{G}$  of subgraphs.
7:   end if
8: end for
9: if  $\bar{G} = \emptyset$  then
10:  add  $G$  to  $\bar{C}$ .
11: end if
12: return  $\bar{G}$ ;
```

---

$G'(V', E')$  and apply Ford & Fulkerson Algorithm [13] to compute the max flow.

2. Once we find nodes  $a, b$  where  $\kappa(a, b) \leq k$ , use the min cut of the flow network as separator  $\bar{S}$ . Remove  $\bar{S}$  and obtain several subgraphs. Add  $\bar{S}$  back to each subgraph and mark  $\bar{S}$  as “separator”. Return the subgraphs for screening.
3. Otherwise,  $G$  is  $k$ -robust and output it as a  $k$ -core.

**EXAMPLE 4.16.** Continue with Example 4.14 and  $k = 2$ . There are four 3-connected v-unions. When we check  $r_1 \in Q_1$  and  $r_6 \in Q_3$ , we find  $\bar{S} = \{r_3, r_4\}$ . We then split  $G_2$  into subgraphs  $\{r_1 - r_4\}$  and  $\{r_3 - r_6\}$ , marking  $r_3$  and  $r_4$  as “separators”.

Now consider graph  $G_1$  in Figure 2 and  $k = 2$ . There are four 3-connected v-unions (actually four v-cliques) and six pairs of adjacent v-unions. For  $Q_1$  and  $Q_2$ , we check nodes  $r_2$  and  $r_4$  and find  $\kappa(r_2, r_4) = 3$ . Similarly we check for every other pair of adjacent v-unions and decide that the graph is 2-robust.  $\square$

**PROPOSITION 4.17.** Let  $p$  be the total number of pairs of adjacent v-unions, and  $g$  be the number of nodes in the input graph. Algorithm SPLIT runs in time  $O(pg^{2.5})$ .  $\square$

**PROOF.** Authors in [9] proves that it takes in time  $O(g^{2.5})$  to compute  $\kappa(a, b)$  for a pair of nodes  $a$  and  $b$  in  $G$ . In the worst case SPLIT needs to compute  $\kappa(a, b)$  for  $p$  pairs of adjacent v-unions. Thus, SPLIT runs in time  $O(pg^{2.5})$ .  $\square$

Recall that if we solve the Max-Flow Problem directly for each pair of sources in the original graph, the complexity is  $O(|\bar{L}|^{4.5})$ , which would be dramatically higher.

### 4.3.3 Full algorithm

We are now ready to present the full algorithm, CORE (Algorithm 3). Initially, it initializes the working queue  $\mathbf{Q}$  with only input  $G$  (Line 1). Each time it pops a subgraph  $G'$  from  $\mathbf{Q}$  and invokes SCREEN (Lines 3-4). If the output of SCREEN is still  $G'$  (so  $G'$  is a  $(k + 1)$ -overlap subgraph) (Line 5), it removes any node with mark “separator” in  $G'$  and puts the new subgraph into the working queue (Line 7), or invokes SPLIT on  $G'$  if there is no separator (Line 9). Subgraphs output by SCREEN and SPLIT are added to the queue for further examination (Lines 10, 13) and identified cores are added to  $\bar{C}$ , the core set. It terminates when  $\mathbf{Q} = \emptyset$ .

---

**Algorithm 3** CORE( $G, k$ )

---

**Input:**  $G$ : Simplified similarity graph, represented by  $\mathbf{C}$  and  $\bar{L}$ .

$k$ : Robustness requirement.

**Output:**  $\bar{C}$  Set of cores in  $G$ .

```
1: Let  $\mathbf{Q} = \{G\}$ ,  $\bar{C} = \emptyset$ ;
2: while  $\mathbf{Q} \neq \emptyset$  do
3:   Pop  $G'$  from  $\mathbf{Q}$ ;
4:   Let  $\bar{P} = \text{SCREEN}(G', k, \bar{C})$ ;
5:   if  $\bar{P} = \{G'\}$  then
6:     if  $G'$  contains “separator” nodes then
7:       Remove separators from  $G'$  and add the result to  $\mathbf{Q}$  if
       it is not empty;
8:     else
9:       Let  $\bar{S} = \text{SPLIT}(G', k, \bar{C})$ ;
10:      add graphs in  $\bar{S}$  to  $\mathbf{Q}$ ;
11:    end if
12:  else
13:    add graphs in  $\bar{P}$  to  $\mathbf{Q}$ ;
14:  end if
15: end while
16: return  $\bar{C}$ ;
```

---

The correctness of algorithm CORE is guaranteed by the following Lemmas.

**LEMMA 4.18.** For each pair of adjacent nodes  $r, r'$  in graph  $G$ , there exists a maximal  $k$ -robust partitioning such that  $r, r'$  are in the same subgraph.  $\square$

**PROOF.** For each pair of adjacent nodes  $r, r'$  in  $G$ , we prove the existence of such a maximal  $k$ -robust partitioning by constructing it.

By definition, adjacent node  $r, r'$  form a v-clique  $C$ . Therefore, there exists a maximal v-clique  $C'$  in  $G$  that contains  $r, r'$ , i.e.,  $C \subseteq C'$ . V-clique  $C'$  can be obtained by keep adding nodes in  $G$  to  $C$  so that each newly-added node is adjacent to each node in current clique until no nodes in  $G$  can be added to  $C'$ . By definition, any v-clique is  $k$ -robust, therefore there exists a maximal  $k$ -robust subgraph  $G'$  in  $G$  such that  $C' \subseteq G'$ . Graph  $G'$  can be obtained by keep adding nodes in  $G$  to  $C'$  so that each newly-added node is adjacent to at least  $k + 1$  nodes in current graph  $G'$  until no nodes in  $G$  can be added to  $G'$ . We remove  $G'$  from  $G$  and take  $G'$  as a subgraph in the desired partitioning.

We repeat the above process to a randomly-selected pair of adjacent nodes in the remaining graph  $G \setminus G'$  until it is empty. The desired partitioning satisfies Condition 1 and 2 of Definition 4.2 because the above process makes sure each subgraph is exclusive and  $k$ -robust; it satisfies Condition 3 of Definition 4.2 because the above process makes sure each subgraph is maximal, which means merging arbitrary number of subgraphs in the partitioning would violate Condition 2.

In summary, the desired partitioning is a maximal  $k$ -robust partitioning. It proves that for each pair of adjacent nodes  $r$  and  $r'$  in graph  $G$ , there exists a maximal  $k$ -robust partitioning such that  $r$  and  $r'$  are in the same subgraph.  $\square$

**LEMMA 4.19.** The set of nodes in a separator  $\bar{S}$  of graph  $G$  does not belong to any  $k$ -core in  $G$ , where  $|\bar{S}| \leq k$ .  $\square$

**PROOF LEMMA 4.19.** Suppose the set  $\bar{S}$  of nodes separate  $G$  into  $m$  disconnected sets  $\bar{X}_i, i \in [1, m], m > 0$ . To prove that each node  $r \in \bar{S}$  does not belong to any  $k$ -core in  $G$ , we prove that for a node  $r' \in G, r' \neq r$ , there exists a maximal  $k$ -robust

partitioning such that  $r$  and  $r'$  are separated. Node  $r'$  falls into the following cases: 1)  $r' \in \bar{X}_i, i \in [1, m]$ ; 2)  $r' \in \bar{S}$ .

Consider Case 1) where  $r' \in \bar{X}_i, i \in [1, m]$ . We construct a maximal  $k$ -robust partitioning of  $G$  where  $r$  and  $r'$  are in different subgraphs. We start with a maximal  $k$ -robust subgraph  $G'$  in  $G$  that contains  $r$  and  $r''$  where  $r''$  is adjacent to  $r$  and in  $\bar{X}_j, j \neq i, j \in [1, m]$ , and find other maximal  $k$ -robust subgraphs as in Lemma 4.18. Since  $\bar{S}$  separates  $\bar{X}_i$  and  $\bar{X}_j$ , maximal  $k$ -robust subgraph  $G'$  that contains  $r$  and  $r''$  does not contain any node in  $\bar{X}_i$ . It proves that there exists a maximal  $k$ -robust partitioning of  $G$  where  $r$  and  $r'$  are not in the same subgraph.

Consider Case 2) where  $r' \in \bar{S}$ . We construct a maximal  $k$ -robust partitioning of  $G$  such that  $r$  and  $r'$  are in different subgraphs. We create two maximal  $k$ -robust subgraphs  $G'$  and  $G''$ , where  $G'$  contains  $r$  and an adjacent node  $r_i \in \bar{X}_i, i \in [1, m]$ ,  $G''$  contains  $r'$  and an adjacent node  $r_j \in \bar{X}_j, j \neq i, j \in [1, m]$ . We create other subgraphs as in Lemma 4.18. Since each path between  $r_i \in \bar{X}_i$  and  $r_j \in \bar{X}_j$  contains at least one node in  $\bar{S}$  and  $|\bar{S}| \leq k$ , graph  $G' \cup G''$  is not  $k$ -robust. Therefore, the created partitioning is a maximal  $k$ -robust partitioning. It proves that there exists a maximal  $k$ -robust partitioning of  $G$  where  $r$  and  $r'$  are not in the same subgraph.

Given the above two cases, we have that any node in separator  $\bar{S}$  of  $G$  does not belong to any  $k$ -core in  $G$ , where  $|\bar{S}| \leq k$ .  $\square$

**THEOREM 4.20.** *Let  $G$  be the input graph and  $q$  be the number of  $(k+1)$ -connected v-unions in  $G$ . Define  $a, p, g, m$ , and  $|\bar{L}|$  as in Proposition 4.11 and 4.17. Algorithm CORE finds correct  $k$ -cores of  $G$  in time  $O(q((m^2 + a)|\bar{L}| + pg^{2.5}))$  and is order independent.*  $\square$

**PROOF.** We first prove that CORE correctly finds  $k$ -cores in  $G$ , that is 1) nodes not returned by CORE do not belong to any  $k$ -core; 2) each subgraph returned by CORE forms a  $k$ -core.

We prove that nodes not returned by CORE do not belong to any  $k$ -core in  $G$ . Nodes not returned by CORE belong to separators of subgraphs in  $G$ . Suppose  $\bar{S}$  is a separator of graph  $G_n \in \mathbf{Q}$  found in either SCREEN or SPLIT phase, where  $G_n \subseteq G, n \geq 0, G_0 = G$ , and  $\bar{S}$  separates  $G_n$  into  $m$  sub-graphs  $\bar{X}_n^i, i \in [1, m], m > 1$ . Graph  $G_n^i \in \mathbf{Q}$  is a subgraph of  $G_n$  such that any node  $r \in \bar{X}_n^j, j \in [1, m], j \neq i$  does not belong to  $G_n^i$ . Nodes removed in  $G_n^i$  by CORE belong to separator  $\bar{S}$  in  $G_n$ . Given Lemma 4.19, such nodes do not belong to any  $k$ -core in  $G_n$  and thus does not belong to any  $k$ -core in  $G$ .

We next prove that each subgraph returned by CORE forms a  $k$ -core in  $G$ . We prove two cases: 1) subgraph  $G'$  in  $G$  forms a  $k$ -core if there exists a separator  $\bar{S}$  that disconnects  $G'$  from  $G$ , where  $|\bar{S}| \leq k$  and  $G' \cup \bar{S}$  and  $G'$  are both  $k$ -robust; 2) if a subgraph is a  $k$ -core in  $G_n^i$ , it is a  $k$ -core in graph  $G_n$ .

We consider Case 1) that subgraph  $G'$  in  $G$  forms a  $k$ -core if there exists a separator  $\bar{S}$  that disconnects  $G'$  from  $G$ , where  $|\bar{S}| \leq k$  and  $G' \cup \bar{S}$  and  $G'$  are both  $k$ -robust. For a pair of nodes  $r_1, r_2$  in  $G'$ , we prove that there exists no maximal  $k$ -robust partitioning where  $r_1$  and  $r_2$  are in different subgraphs. Suppose such a partitioning exists, and  $G_1, G_2$  are subgraphs containing  $r_1, r_2$  respectively. Since  $G_1, G_2 \subseteq G' \cup \bar{S}$ , we have that  $G_1 \cup G_2$  is  $k$ -robust, it violates the fact that the result of merging any two subgraphs in a maximal  $k$ -robust partitioning is not  $k$ -robust. Therefore, there exists no maximal  $k$ -robust partitioning where  $r_1$  and  $r_2$  are in different subgraphs. It proves that  $G'$  is a  $k$ -core in  $G$ .

We next consider Case 2) that if a subgraph  $G'$  is a  $k$ -core in  $G_n^i$ , it is a  $k$ -core in graph  $G_n$ . We prove that a pair of nodes  $r_1, r_2 \in G'$  belong to the same subgraph of all maximal  $k$ -robust partitioning in  $G_n$ . Suppose there exists such a partitioning of  $G_n$  where  $r_1 \in G_1, r_2 \in G_2$ . Since  $G_n^i \subseteq \bar{X}_n^i \cup \bar{S}$ , we have  $G_1, G_2 \subseteq$

**Table 4: Step-by-step core identification in Example 4.21.**

Input	Method	Output
$G_2$	SCREEN	$G_2$
$G_2$	SPLIT	$G_2^1 = \{r_1 - r_4\}, G_2^2 = \{r_3 - r_6\}$
$G_2^1$	SCREEN	$G_2^3 = \{r_3\}, G_2^4 = \{r_4\}$
$G_2^2$	SCREEN	$G_2^5 = \{r_3\}, G_2^6 = \{r_4\}$
$G_2^3$	SCREEN	-
$G_2^4$	SCREEN	-
$G$	SCREEN	$G^1 = \{r_1/\dots/7\}, G^2 = \{r_{11}, r_{12}, r_{14}/15\}, G^3 = \{r_{16}/17/18\}$
$G^1$	SCREEN	Core $\{r_1 - r_7\}$
$G^2$	SCREEN	$G^4 = \{r_{11}\}, G^5 = \{r_{14}/15\}$
$G^3$	SCREEN	Core $\{r_{16} - r_{18}\}$
$G^4$	SCREEN	-
$G^5$	SCREEN	Core $\{r_{14} - r_{15}\}$

$G_n^i$ , otherwise  $G_1, G_2$  are not  $k$ -robust. Since  $r_1, r_2$  belong to the same  $k$ -core in  $G_n^i$ , we have  $G_1 = G_2$ . It proves that if  $G'$  is a  $k$ -core in  $G_n^i$ , it is a  $k$ -core in  $G_n$ .

The above two cases prove that each subgraph returned by CORE forms a  $k$ -core in  $G$ . In summary, nodes not returned by CORE do not belong to any  $k$ -core, and each subgraph returned by CORE forms a  $k$ -core in  $G$ . Thus, CORE correctly finds all  $k$ -cores in  $G$ . It further proves that the result of CORE is independent from the order in which we find and remove separators of graphs in  $\mathbf{Q}$ .

We now analyze the time complexity of CORE. For each  $(k+1)$ -connected v-unions in  $G$ , it takes in time  $O(m^2 + a)|\bar{L}|$  to proceed SCREEN phase and in time  $O(pg^{2.5})$  to proceed SPLIT phase. In total there are  $q$  v-unions in  $G$ , thus the algorithm takes in time  $O(q((m^2 + a)|\bar{L}| + pg^{2.5}))$ .  $\square$

**EXAMPLE 4.21.** *First, consider graph  $G_2$  in Figure 2 and  $k = 2$ . Table 4 shows the step-by-step core identification process. It passes screening and is the input for SPLIT. SPLIT then splits it into  $G_2^1$  and  $G_2^2$ , where  $r_3$  and  $r_4$  are marked as “separators”. SCREEN further splits each of them into  $\{r_3\}$  and  $\{r_4\}$ , both discarded as each represents a single node (and is a separator). So CORE does not output any core.*

*Next, consider the motivating example, with the input shown in Table 3 and  $k = 1$ . Originally,  $\mathbf{Q} = \{G\}$ . After invoking SCREEN on  $G$ , we obtain three subgraphs  $G^1, G^2$ , and  $G^3$ . SCREEN outputs  $G^1$  and  $G^3$  as 1-cores since each contains a single node that represents multiple records. It further splits  $G^2$  into two single-node graphs  $G^4$  and  $G^5$ , and outputs the latter as a 1-core. Note that if we remove the 1-robustness requirement, we would merge  $r_{11} - r_{18}$  to the same core and get false positives.*  $\square$

**Case study:** On the data set with 18M records, our core-identification algorithm finished in 2.2 minutes. SCREEN was invoked 114K times and took 2 minutes (91%) in total. Except the original graph, an input contains at most 39.3K nodes; for 97% inputs there are fewer than 10 nodes and running SCREEN was very fast. SPLIT was invoked only 26 times; an input contains at most 65 nodes (13 v-unions) and on average 7.8 (2.7 v-unions). Recall that the simplified inverted index contains 1.5M entries, so SCREEN reduced the size of the input to SPLIT by 4 orders of magnitude.

## 5. GROUP LINKAGE

The second stage clusters the cores and the remaining records, which we call *satellites*, into groups. To avoid merging records based only on weak evidence, we require that a cluster cannot contain more than one satellite but no core. Comparing with clustering in traditional record linkage, our algorithm differs in three aspects. First, in addition to weighting each attribute, we weight the values

according to their popularity within a group such that similarity on primary values (strong evidence) is rewarded more. Second, we treat all values for dominant-value attributes as a whole, we are tolerant to differences on local values from different entities in the same group. Third, we distinguish weights for distinct values and non-distinct values such that similarity on distinct values is rewarded more. This section first describes the objective function for clustering (Section 5.1) and then proposes a greedy algorithm for clustering (Section 5.2).

## 5.1 Objective function

**SV-index:** Ideally, we wish that each cluster is *cohesive* (each element, being a core or a satellite, is close to other elements in the same cluster) and different clusters are *distinct* (each element is fairly different from those in other clusters). Since records in the same group may have fairly different local values, we adopt *Silhouette Validation Index (SV-index)* [25] as the objective function as it is more tolerant to diversity within a cluster. Given a clustering  $\mathcal{C}$  of elements  $\mathbf{E}$ , the SV-index of  $\mathcal{C}$  is defined as follows.

$$S(\mathcal{C}) = \text{Avg}_{e \in \mathbf{E}} S(e); \quad (1)$$

$$S(e) = \frac{a(e) - b(e) + \alpha}{\max\{a(e), b(e)\} + \beta}. \quad (2)$$

Here,  $a(e) \in [0, 1]$  denotes the similarity between element  $e$  and its own cluster,  $b(e) \in [0, 1]$  denotes the maximum similarity between  $e$  and another cluster,  $\beta > \alpha > 0$  are small numbers to keep  $S(e)$  finite and non-zero (we discuss in Section 6 how we set the parameters). A nice property of  $S(e)$  is that it falls in  $[-1, 1]$ , where a value close to 1 indicates that  $e$  is in an appropriate cluster, a value close to  $-1$  indicates that  $e$  is mis-classified, and a value close to 0 while  $a(e)$  is not too small indicates that  $e$  is equally similar to two clusters that should possibly be merged. Accordingly, we wish to obtain a clustering with the maximum SV-index. We next describe how we compare an element with a cluster.

**Similarity computation:** We consider that an element  $e$  is similar to a cluster  $Cl$  if they have highly similar values on common-value attributes (e.g., `name`), share at least one *primary* value (we explain “primary” later) on dominant-value attributes (e.g., `phone`, `URL`); in addition, our confidence is higher if they also share values on multi-value attributes (e.g., `category`). Following previous work on handling multi-value attributes [7, 21], we compute the similarity  $\text{sim}(e, Cl)$  as follows.

$$\text{sim}(e, Cl) = \min\{1, \text{sim}_s(e, Cl) + \tau w_m \text{sim}_{\text{multi}}(e, Cl)\}; \quad (3)$$

$$\text{sim}_s(e, Cl) = \frac{w_c \text{sim}_{\text{com}}(e, Cl) + w_o \text{sim}_{\text{dom}}(e, Cl)}{w_c + w_o}; \quad (4)$$

$$\tau = \begin{cases} 0 & \text{if } \text{sim}_s(e, Cl) < \theta_{th}, \\ 1 & \text{otherwise.} \end{cases} \quad (5)$$

Here,  $\text{sim}_{\text{com}}$ ,  $\text{sim}_{\text{dom}}$ , and  $\text{sim}_{\text{multi}}$  denote the similarity for common-, dominant-, and multi-attributes respectively. We take the weighted sum of  $\text{sim}_{\text{com}}$  and  $\text{sim}_{\text{dom}}$  as strong indicator of  $e$  belonging to  $Cl$  (measured by  $\text{sim}_s(e, Cl)$ ), and only reward weak indicator  $\text{sim}_{\text{multi}}$  if  $\text{sim}_s(e, Cl)$  is above a pre-defined threshold  $\theta_{th}$ ; the similarity is at most 1. Weights  $0 < w_c, w_o, w_m < 1$  indicate how much we reward value similarity or penalize value difference; we learn the weights from sampled data. We next highlight how we leverage strong evidence from cores and meanwhile remain tolerant to other different values in similarity computation.

First, we identify *primary values* (strong evidence) as popular values within a cluster. When we maintain the signature for a core or a cluster, we keep all values of an attribute and assign a high *weight* to a popular value. Specifically, let  $\bar{R}$  be a set of records.

Consider value  $v$  and let  $\bar{R}(v) \subseteq \bar{R}$  denote the records in  $\bar{R}$  that contain  $v$ . The weight of  $v$  is computed by  $w(v) = \frac{|\bar{R}(v)|}{|\bar{R}|}$ .

**EXAMPLE 5.1.** Consider phone for core  $Cr_1 = \{r_1 - r_7\}$  in Table 2. There are 7 business listings in  $Cr_1$ , 5 providing 808 ( $r_1 - r_5$ ), one providing 101 ( $r_6$ ), and one providing 102 ( $r_7$ ). Thus, the weight of 808 is  $\frac{5}{7} = .71$  and the weight for 101 and 102 is  $\frac{1}{7} = .14$ , showing that 808 is the primary phone for  $Cr_1$ .  $\square$

Second, when we compute  $\text{sim}_{\text{dom}}(e, Cl)$ , we consider all the dominant-value attributes together, rewarding sharing primary values (values with a high weight) but not penalizing different values unless there is no shared value. Specifically, if the primary value of an element is the same as that of a cluster, we consider them having probability  $p$  to be in the same group. Since we use weights to measure whether the value is primary and allow slight difference on values, with a value  $v$  from  $e$  and  $v'$  from  $Cl$ , the probability becomes  $p \cdot w_e(v) \cdot w_{Cl}(v') \cdot s(v, v')$ , where  $w_e(v)$  measures the weight of  $v$  in  $e$ ,  $w_{Cl}(v')$  measures the weight of  $v'$  in  $Cl$ , and  $s(v, v')$  measures the similarity between  $v$  and  $v'$ . We compute  $\text{sim}_{\text{dom}}(e, Cl)$  as the probability that they belong to the same group given several shared values as follows.

$$\text{sim}_{\text{dom}}(e, Cl) = 1 - \prod_{v \in e, v' \in ch} (1 - p \cdot w_e(v) \cdot w_{Cl}(v') \cdot s(v, v')). \quad (6)$$

When there is no shared primary value,  $\text{sim}_{\text{dom}}$  can be close to 0; once there is one such value,  $\text{sim}_{\text{dom}}$  can be significantly increased, since we typically set a large  $p$ .

**EXAMPLE 5.2.** Consider element  $e = r_8$  and cluster  $Cl_1 = \{r_1 - r_7\}$  in Example 1.1. Assume  $p = .9$ . Element  $e$  and  $Cl_1$  share the primary email domain, with weight 1 and  $\frac{5}{7} = .71$  respectively, but have different phone numbers (assuming similarity of 0). We compute  $\text{sim}_{\text{dom}}(e, Cl_1) = 1 - (1 - .9 \cdot 1 \cdot .71 \cdot 1) \cdot (1 - 0) \cdot (1 - 0) \cdot (1 - 0) = .639$ ; essentially, we do not penalize the difference in phone numbers. Note however if *homedepot* appeared only once so was not a primary value, its weight would be .14 and accordingly  $\text{sim}_{\text{dom}}(e, Cl_1) = .126$ , indicating a much lower similarity.  $\square$

Third, when we learn weights, we learn one set of weights for distinct values (appearing in only one cluster) and one set for non-distinct values, such that distinct values, which can be considered as stronger evidence, typically contribute more to the final similarity. In Example 1.1, sharing “*Home Depot, The*” would serve as stronger evidence than sharing *Taco Casa* for group similarity.

## 5.2 Clustering algorithm

In most cases, clustering is intractable [14, 26]. We maximize the SV-index in a greedy fashion. Our algorithm starts with an initial clustering and then iteratively examines if we can improve the current clustering (increase SV-index) by merging clusters or moving elements between clusters. According to the definition of SV-index, in both initialization and adjusting, we always assign an element to the cluster with which it has the highest similarity.

**Initialization:** Initially, we (1) assign each core to its own cluster and (2) assign a satellite  $r$  to the cluster with the highest similarity if the similarity is above threshold  $\theta_{ini}$  and create a new cluster for  $r$  otherwise. We update the signature of each core along the way. Note that initialization is sensitive in the order we consider the records. Although designing an algorithm independent of the ordering is possible, such an algorithm is more expensive and our experiments show that the iterative adjusting can smooth out the difference.




Figure 4: Clustering of  $r_{11} - r_{20}$  in Table 2.

Table 5: Element-cluster similarity and SV-index for clusterings in Figure 4. Similarity between an element and its own cluster is in bold and the second-to-highest similarity is in italic. Low  $S(e)$  scores are in italic.

	Cl <sub>2</sub>	Cl <sub>3</sub>	Cl <sub>4</sub>	Cl <sub>5</sub>	Cl <sub>6</sub>	$S(e)$
Cr <sub>2</sub>	<b>.9</b>	.5	.5	.5	.5	.44
Cr <sub>3</sub>	.6	<b>1</b>	.5	.5	.5	.4
$r_{11}$	.7	.5	<b>1</b>	.5	.5	.3
$r_{12}$	<b>.99</b>	.5	.95	.5	.5	.05
$r_{13}$	<b>1</b>	.9	.95	.5	.5	.05
$r_{19}$	.5	.5	.5	<b>1</b>	.5	.5
$r_{20}$	.5	.5	.5	.5	<b>1</b>	.5

(a) Cluster  $C_a$ .

	Cl <sub>2</sub>	Cl <sub>3</sub>	Cl <sub>5</sub>	Cl <sub>6</sub>	$S(r)$
Cr <sub>2</sub>	<b>.87</b>	.5	.5	.5	.43
Cr <sub>3</sub>	.58	<b>1</b>	.5	.5	.42
$r_{11}$	<b>.79</b>	.5	.5	.5	.37
$r_{12}$	<b>.96</b>	.5	.5	.5	.48
$r_{13}$	<b>.97</b>	.9	.5	.5	.07
$r_{19}$	.5	.5	<b>1</b>	.5	.5
$r_{20}$	.5	.5	.5	<b>1</b>	.5

(b) Cluster  $C_b$ .

EXAMPLE 5.3. Continue with the motivating example in Table 2. First, consider records  $r_1 - r_{10}$ , where  $Cr_1 = \{r_1 - r_7\}$  is a core. We first create a cluster  $Cl_1$  for  $Cr_1$ . We then merge records  $r_8 - r_{10}$  to  $Cl_1$  one by one, as they share similar names, and either primary phone number or primary URL.

Now consider records  $r_{11} - r_{20}$ ; recall that there are 2 cores and 5 satellites after core identification. Figure 4 shows the initialization result  $C_a$ . Initially we create two clusters  $Cl_2, Cl_3$  for cores  $Cr_2, Cr_3$ . Records  $r_{11}, r_{19} - r_{20}$  do not share any primary value on dominant-value attributes with  $Cl_2$  or  $Cl_3$ , so have a low similarity with them; we create a new cluster for each of them. Records  $r_{12}$  and  $r_{13}$  share the primary phone with  $Cr_2$  so have a high similarity; we link them to  $Cl_2$ .  $\square$

**Cluster adjusting:** Although we always assign an element  $e$  to the cluster with the highest similarity so  $S(e) > 0$ , the result clustering may still be improved by merging some clusters or moving a subset of elements from one cluster to another. Recall that when  $S(e)$  is close to 0 and  $a(e)$  is not too small, it indicates that a pair of clusters might be similar and is a candidate for merging. Thus, in cluster adjusting, we find such candidate pairs, iteratively adjust them by merging them or moving a subset of elements between them, and choose the new clustering if it increases the SV-index.

We first describe how we find candidate pairs. Consider element  $e$  and assume it is closest to clusters  $Cl$  and  $Cl'$ . If  $S(e) \leq \theta_s$ , where  $\theta_s$  is a threshold for considering merging, we call it a border element of  $Cl$  and  $Cl'$  and consider  $(Cl, Cl')$  as a candidate pair. We rank the candidates according to (1) how many border elements they have and (2) for each border element  $e$ , how close  $S(e)$  is to 0. Accordingly, we define the *benefit* of merging  $Cl$  and  $Cl'$  as  $b(Cl, Cl') = \sum_{e \text{ is a border of } Cl \text{ and } Cl'} (1 - S(e))$ , and rank the candidate pairs in decreasing order of the benefit.

We next describe how we re-cluster elements in a candidate pair  $(Cl, Cl')$ . We adjust by merging the two clusters, or moving the border elements between the clusters, or moving out the border el-




Figure 5: Reclustering plans for  $Cl_1$  and  $Cl_2$ .

ements and merging them. Figure 5 shows the four re-clustering plans for a candidate pair. Among them, we consider those that are valid (*i.e.*, a cluster cannot contain more than one satellite but no core) and choose the one with the highest SV-index. When we compute SV-index, we consider only elements in  $Cl$  and  $Cl'$  and those that are second-to-closest to  $Cl$  or  $Cl'$  (their  $a(e)$  or  $b(e)$  can be changed) such that we can reduce the computation cost. After the adjusting, we need to re-compute  $S(e)$  for these elements and update the candidate-pair list accordingly.

EXAMPLE 5.4. Consider adjusting cluster  $C_a$  in Figure 4. Table 5(a) shows similarity of each element-cluster pair and SV-index of each element. Thus, the SV-index is .32.

Suppose  $\theta_s = .3$ . Then,  $r_{11} - r_{13}$  are border elements of  $Cl_2$  and  $Cl_4$ , where  $b(Cl_2, Cl_4) = .7 + .95 = 2.6$  (there is a single candidate so we do not need to compare the benefit). For the candidate, we have two re-clustering plans,  $\{\{r_{11} - r_{13}, Cr_2\}\}$ ,  $\{\{r_{11} - r_{13}\}, \{Cr_2\}\}$ , while the latter is invalid. For the former ( $C_b$  in Figure 4), we need to update  $S(e)$  for every element and the new SV-index is .4 (Table 5(b)), higher than the original one.  $\square$

The full clustering algorithm CLUSTER (details in Algorithm 4) goes as follows.

1. Initialize a clustering  $\mathcal{C}$  and a list  $Que$  of candidate pairs ranked in decreasing order of merging benefit. (Lines 1-2).
2. For each candidate pair  $(Cl, Cl')$  in  $Que$  do the following.
  - (a) Examine each valid adjusting plan and compute SV-index for it, and choose the one with the highest SV-index. (Line 4).
  - (b) Change the clustering if the new plan has a higher SV-index than the original clustering. Recompute  $S(e)$  for each relevant element  $e$  and move  $e$  to a new cluster if appropriate. Update  $Que$  accordingly. (Lines 6-16).
3. Repeat Step 2 until  $Que = \emptyset$ .

PROPOSITION 5.5. Let  $l$  be the number of distinct candidate pairs ever in  $Que$  and  $|\mathbf{E}|$  be the number of input elements. Algorithm CLUSTER takes time  $O(l \cdot |\mathbf{E}|^2)$ .  $\square$

PROOF. It takes time  $O(|\mathbf{E}|^2)$  to initialize clustering  $\mathcal{C}$  and list  $Que$ . It takes  $|\mathbf{E}|^2$  to check each distinct candidate pair in  $Que$ , where it takes  $O(|\mathbf{E}|)$  to examine all valid clustering plans and select the one with highest SV-index (Step 2(a)), and it takes  $O(|\mathbf{E}|^2)$  to recompute SV-index for all relevant elements and update  $Que$  (Step 2(b)). In total there are  $l$  distinct candidate pairs ever in  $Que$ , thus CLUSTER takes time  $O(l \cdot |\mathbf{E}|^2)$ .  $\square$

Note that we first block records according to name similarity and take each block as an input, so typically  $|\mathbf{E}|$  is quite small. Also, in practice we need to consider only a few candidate pairs for adjusting in each input, so  $l$  is also small.

**Algorithm 4** CLUSTER( $\mathbf{E}, \theta_s$ )

---

**Input:**  $\mathbf{E}$ : A set of cores and satellites for clustering.  
 $\theta_s$ : Pre-defined threshold for considering merging.  
**Output:**  $\mathcal{C}$ : A clustering of elements in  $\mathbf{E}$ .

---

```

1: Initialize  $\mathcal{C}$  according to  $\mathbf{E}$ ;
2: Compute  $S(\mathcal{C})$  and generate a list  $Que$  of candidate pairs;
3: for each candidate pair  $(Cl, Cl') \in Que$  do
4:   compute SV-index for its valid re-clustering plans and
     choose the clustering  $\mathcal{C}_{max}$  with the highest SV-index;
5:   if  $S(\mathcal{C}) < S(\mathcal{C}_{max})$  then
6:     let  $\mathcal{C} = \mathcal{C}_{max}$ ,  $change = true$ ;
7:     while  $change$  do
8:        $change = false$ ;
9:       for each relevant element  $e$  do
10:        recompute  $S(e)$ ;
11:        When appropriate, move  $e$  to a new cluster and set
          $change = true$ ;
12:        if  $S(e) < \theta_s$  in the previous or current  $\mathcal{C}$  then
13:          update the merging benefit of the related candi-
           date pair and add it to  $Que$  or remove it from  $Que$ 
           when appropriate;
14:        end if
15:      end for
16:    end while
17:  end if
18: end for
19: return  $\mathcal{C}$ ;

```

---

**Table 6: Statistics of the experimental data sets.**

	#Records	#Groups (size > 1)	Group size	#Singletons (size = 1)
<i>Random</i>	2062	30	[2, 308]	503
<i>AI</i>	2446	1	2446	0
<i>UB</i>	322	9	[2, 275]	5
<i>FBIIns</i>	1149	14	[33, 269]	0
<i>SIGMOD</i>	590	71	[2, 41]	162

EXAMPLE 5.6. Continue with Example 5.4 and consider adjusting  $\mathcal{C}_b$ . Now there is one candidate pair  $(Cl_2, Cl_3)$ , with border  $r_{13}$ . We consider clusterings  $\mathcal{C}_c$  and  $\mathcal{C}_d$ . Since  $S(\mathcal{C}_c) = .37 < .40$  and  $S(\mathcal{C}_d) = .32 < .40$ , we keep  $\mathcal{C}_b$  and return it as the result. We do not merge records  $Cl_2 = \{r_{11} - r_{15}\}$  with  $Cl_3 = \{r_{16} - r_{18}\}$ , because they share neither phone nor the primary URL. CLUSTER returns the correct chains.  $\square$

## 6. EXPERIMENTAL EVALUATION

This section describes experimental results on two real-world data sets, showing high scalability of our techniques, and advantages of our algorithm over rule-based or traditional machine-learning methods on accuracy.

### 6.1 Experiment settings

**Data and gold standard:** We experimented on two real-world data sets. *Biz* contains 18M US business listings and each listing has attributes name, phone, URL, location and category; we decide which listings belong to the same business chain. *SIGMOD* contains records about 590 attendees of SIGMOD’98 and each record has attributes name, affiliation, address, phone, fax and email; we decide which attendees belong to the same institute.

We experimented on the whole *Biz* data set to study scalability of our techniques. We evaluated accuracy of our techniques on five subsets of data. The first four are from *Biz*. (1) *Random* contains 2062 listings from *Biz*, where 1559 belong to 30 randomly selected business chains, and 503 do not belong to any chain; among the 503

listings, 86 are highly similar in name to listings in the business chains and the rest are randomly selected. (2) *AI* contains 2446 listings for the same business chain *Allstate Insurance*. These listings have the same name, but 1499 provide URL “*allstate.com*”, 854 provide another URL “*allstateagencies.com*”, while 130 provide both, and 227 listings do not provide any value for phone or URL. (3) *UB* contains 322 listings with exactly the same name *Union Bank* and highly similar category values; 317 of them belong to 9 different chains while 5 do not belong to any chain. (4) *FBIIns data set* contains 1149 listings with similar names and highly similar category values; they belong to 14 different chains. Among the listings, 708 provide the same wrong name *Texas Farm Bureau Insurance* and meanwhile provide a wrong URL *farmbureauinsurance-mi.com*. Among these four subsets, the latter three are hard cases; for each data set, we manually verified all the chains by checking store locations provided by the business-chain websites and used it as the gold standard. The last “subset” is actually the whole *SIGMOD* data set. It has very few wrong values, but the same affiliation can be represented in various ways and some affiliation names can be very similar (e.g., *UCSC* vs. *UCSD*). We manually identified 71 institutes that have multiple attendees and there are 162 attendees who do not belong to these institutes. Table 6 shows statistics of the five subsets.

**Measure:** We considered each group as a cluster and compared pairwise linking decisions with the gold standard. We measured the quality of the results by *precision* ( $P$ ), *recall* ( $R$ ), and *F-measure* ( $F$ ). If we denote the set of true-positive pairs by  $TP$ , the set of false-positive pairs by  $FP$ , and the set of false-negative pairs by  $FN$ , then,  $P = \frac{|TP|}{|TP|+|FP|}$ ,  $R = \frac{|TP|}{|TP|+|FN|}$ ,  $F = \frac{2PR}{P+R}$ . In addition, we reported execution time.

**Implementation:** We implemented the technique we proposed in this paper, and call it GROUP. In core generation, for *Biz* we considered two records are similar if (1) their name similarity is above .95; and (2) they share at least one phone or URL domain name. For *SIGMOD* we require (1) affiliation similarity is above .95; and (2) they share at least one of phone prefix (3-digit), fax prefix (3-digit), email server, or the addresses have a similarity above .9. We required 2-robustness for cores. In clustering, (1) for blocking, we put records whose name similarity is above .8 in the same block; (2) for similarity computation, we computed string similarity by Jaro-Winkler distance [5], we set  $\alpha = .01$ ,  $\beta = .02$ ,  $\theta_{th} = .6$ ,  $p = .8$ , and we learned other weights from 1000 records randomly selected from *Random* data for *Biz*, and 300 records randomly selected from *SIGMOD*. We discuss later the effect of these choices.

For comparison, we also implemented the following baselines:

- SAMENAME groups *Biz* records with highly similar names and groups *SIGMOD* records with highly similar affiliations (similarity above .95);
- CONNECTEDGRAPH generates the similarity graph as GROUP but considers each connected subgraph as a group;
- One-stage machine-learning linkage methods include PARTITION, CENTER and MERGE [16]; each method computes record similarity by Eq.(3) with learned weights.
- Two-stage method YOSHIDA [30] generates cores by agglomerative clustering with threshold .9 in the first stage, uses TF/IDF weights for features and applies linear algebra to assign each record to a group in the second stage.

We implemented the algorithms in Java. We used a Linux machine with Intel Xeon X5550 processor (2.66GHz, cache 8MB, 6.4GT/s QPI). We used MySQL to store the data sets and stored the index as a database table. Note that after blocking, we can fit

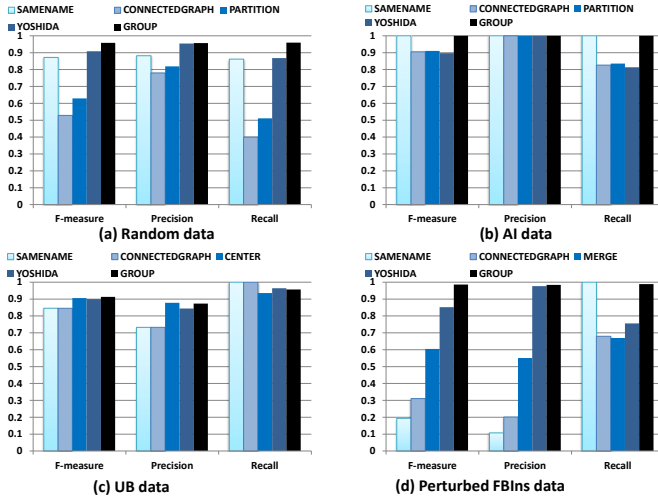


Figure 6: Overall results on *Biz* data set.




Figure 7: Results on *SIGMOD* data.

each block of nodes or elements into main memory, which is typically the case with a good blocking strategy.

## 6.2 Evaluating effectiveness

We first evaluate effectiveness of our algorithms. Figure 6 and Figure 7(a) compare GROUP with the baseline methods, where for the three one-stage linkage methods we plot only the best results. On *FBIns*, all methods put all records in the same chain because a large number (708) of listings have both a wrong name and a wrong URL. We manually perturbed the data as follows: (1) among the 708 listings with wrong URLs, 408 provide a single (wrong) URL and we fixed it; (2) for all records we set `name` to “*Farm Bureau Insurance*”, so removed hints from business names. Even after perturbing, this data set remains the hardest and we use it hereafter instead of the original one for other experiments.

We have the following observations. (1) GROUP obtains the highest F-measure (above .9) on each data set. It has the highest precision most of the time as it applies core identification and leverages the strong evidence collected from resulting cores. It also has a very high recall (mostly above .95) on each subset because the clustering phase is tolerant to diversity of values within chains. (2) The F-measure of SAMENAME is 7-80% lower than GROUP. It can have false positives when listings of highly similar names belong to different chains and can also have false negatives when some listings in a chain have fairly different names from other listings. It only performs well in *AI*, where it happens that all listings have the same name and belong to the same chain. (3) The F-measure of CONNECTEDGRAPH is 2-39.4% lower than SAMENAME. It requires in addition sharing at least one value for dominant-value attributes. As a result, it has a lower recall than SAMENAME; it has fewer false positives than SAMENAME, but because it has fewer true positives, its precision can appear to be lower too. (4) The highest F-measure of one-stage linkage methods is 1-94.7% higher




Figure 8: Contribution of different components on *Biz*.

than CONNECTEDGRAPH. As they require high record similarity, it has similar number of false positives to CONNECTEDGRAPH but often has much more true positives; thus, it often has a higher recall and also a higher precision. However, the highest F-measure is still 1-38.7% lower than GROUP. (5) YOSHIDA has comparable precision to GROUP since its first stage is conservative too, which makes it often improve over the best of one-stage linkage methods on *Biz* dataset where reducing false positives is a big challenge; on the other hand, its first stage is often too conservative (requiring high record similarity) so the recall is 10-34.6% lower than GROUP, which also makes it perform worse than one-stage linkage methods on *Sigmod* dataset where reducing false negatives is challenging.

**Contribution of different components:** We compared GROUP with (1) CORE, which applies Algorithm COREIDENTIFICATION but does not apply clustering, and (2) CLUSTER, which considers each individual record as a core and applies Algorithm CLUSTER (in the spirit of [20, 28]). Figure 8 and Figure 7(b) show the results. First, we observe that CORE improves over one-stage linkage methods on precision by .1-78.6% but has a lower recall (1.5-34.3% lower) most of the time, because it sets a high requirement for merging records into groups. Note however that its goal is indeed to obtain a high precision such that the strong evidence collected from the cores are trustworthy for the clustering phase. Second, CLUSTER often has higher precision (by 1.6-77.3%) but lower recall (by 2.5-32.2%) than the best one-stage linkage methods; their F-measures are comparable on each data set. On some data sets (*Random*, *FBIns*) it can obtain an even higher precision than CORE, because CORE can make mistakes when too many records have erroneous values, but CLUSTER may avoid some of these mistakes by considering also similarity on *state* and *category*. However, applying clustering on the results of CLUSTER would not change the results, but applying clustering on the results of CORE can obtain a much higher F-measure, especially a higher recall (98% higher than CLUSTER on *Random*). This is because the result of CLUSTER lacks the strong evidence collected from high-quality cores so the final results would be less tolerant to diversity of values, showing the importance of core identification. Finally, we observe that GROUP obtains the best results in most of the data sets.

We next evaluate various choices in the two stages. Unless specified otherwise, we observed similar patterns on each data set from *Biz* and *Sigmod*, and report the results on *Random* or perturbed *FBIns* data, whichever has more distinguishable results.

### 6.2.1 Core identification






Figure 9: Core identification on perturbed *FBIns* data.




Figure 10: Effect of graph generation on *Random* data.

**Core identification:** We first compared three core-generation strategies: CORE iteratively invokes SCREEN and SPLIT, ONLYSCREEN only iteratively invokes SCREEN, and YOSHIDAI generates cores by agglomerative clustering [30]. Recall that by default we apply CORE. Figure 9 compares them on the perturbed *FBIns* data. First, we observe similar results of ONLYSCREEN and CORE on all data sets since most inputs to SPLIT pass the  $k$ -robustness test. Thus, although SCREEN in itself cannot guarantee soundness of the resulting cores ( $k$ -robustness), it already does well in practice. Second, YOSHIDAI has lower recall in both core and clustering results, since it has stricter criteria in core generation.

**Graph generation:** We compared three edge-adding strategies for similarity graphs: SIM takes weighted similarity on each attribute except location and requires a similarity of over .8; TWODOM requires sharing name and at least two values on dominant-value attributes; ONEDOM requires sharing name and one value on dominant-value attributes. Recall that by default we applied ONEDOM. Figure 10 compares these three strategies. We observe that (1) SIM requires similar records so has a high precision, with a big sacrifice on recall for the cores (0.00025); as a result, the F-measure of the chains is very low (.59); (2) TWODOM has the highest requirements and so even lower recall than SIM for the cores (.00002), and in turn it has the lowest F-measure for the chains (.52). This shows that only requiring high precision for cores with big sacrifice on recall can also lead to low F-measure for the chains.

We also varied the similarity requirement for names and observed very similar results (varying by .04%) when we varied the threshold from .8 to .95.

**Robustness requirement:** We next studied how the robustness requirement can affect the results (Figure 11). We have three observations. (1) When  $k = 0$ , we essentially take every connected subgraph as a core, so the generated cores can have a much lower precision; those false positives cause both a low precision and a low recall for the resulting chains because we do not collect high-quality strong evidence. (2) When we vary  $k$  from 1 to 4, the number of false positives decreases while that of false negatives increases for the cores, and the F-measure of the chains increases but only very slightly. (3) When we continue increasing  $k$ , the results of cores and clusters remain stable. This is because setting




Figure 11: Effect of robustness requirement on *Random* data.




Figure 12: Clustering strategies on *Random* data.

$k=4$  already splits the graph into subgraphs, each containing a single  $v$ -clique, so further increasing  $k$  would not change the cores. This shows that considering  $k$ -robustness is important, but  $k$  does not need to be too high.

## 6.2.2 Clustering

**Clustering strategy:** We first compared our clustering algorithm with two algorithms proposed for the second stage of two-stage clustering: LIUII [22] iteratively applies majority voting to assign each record to a cluster and collects a set of representative features for each cluster using a threshold (we set it to 5, which leads to the best results); YOSHIDAI [30] is the second stage of YOSHIDA. Figure 12(a) compares their results. We observe that our clustering method improves the recall by 39% over LIUII and by 11% over YOSHIDAI. LIUII may filter strong evidence by the threshold; YOSHIDAI cannot handle records whose dominant-value attributes have null values well.

We also compared four clustering algorithms: GREEDYINITIAL performs only initialization as we described in Section 5; EXHAUSTIVEINITIAL also performs only initialization, but by iteratively conducting matching and merging until no record can be merged to any core; CLUSTERWGREEDY applies cluster adjusting on the results of GREEDYINITIAL, and CLUSTERWEXHAUSTIVE applies cluster adjusting on the results of EXHAUSTIVEINITIAL. Recall that by default we apply CLUSTERWGREEDY. Figure 12(b) compares their results. We observe that (1) applying cluster adjusting can improve the F-measure a lot (by 8.6%), and (2) exhaustive initialization does not significantly improve over greedy initialization, if at all. This shows effectiveness of the current algorithm CLUSTER.

**Value weight:** We then compared the results with and without setting popularity weights for values. Figure 13 compares the results with and without setting popularity weights on perturbed *FBIns* data. We observe that setting the popularity weight helps distinguish primary values from unpopular values, thus can improve the precision. Indeed, on perturbed *FBIns* data it improves the precision from .11 to .98, and improves the F-measure by 403%.

**Attribute weight:** We next considered our weight learning strategy. We first compared SEPARATEDDOMINANT, which learns separated weights for different dominant-value attributes, and UNITED-




Figure 13: Value weights on perturbed *FBIns* data.




Figure 14: Dominant-value attributes on *Random*.




Figure 15: Distinct values on *Random* data.




Figure 16: Attribute weights on *Random* data.

DOMINANT (our default), which considers all such attributes as a whole and learns one single weight for them. Figure 14 shows that on *Random* the latter improves over the former by 95.4% on recall and obtains slightly higher precision, because it penalizes only if neither phone nor URL is shared and so is more tolerant to different values for dominant-value attributes. This shows importance of being tolerant to value variety on dominant-value attributes.

Next, we compared SINGLEWEIGHT, which learns a single weight for each attribute, and DOUBLEWEIGHT (our default), which learns different weights for distinct values and non-distinct values for each attribute. Figure 15 shows that DOUBLEWEIGHT significantly improves the recall (by 94% on *Random*) since it rewards sharing of distinct values, and so can link some satellite records with null values on dominant-value attributes to the chains they should belong to. This shows importance of distinguishing distinct and non-distinct values.

We also compared three weight-setting strategies: (1) 3EQUAL considers common-value attributes, dominant-value attributes, and multi-value attributes, and sets the same weight for each of them; (2) 2EQUAL sets equal weight of .5 for common-value attributes and dominant-value attributes, and weight of .1 for each multi-value attribute; (3) LEARNED applies weights learned from labeled data. Recall that by default we applied LEARNED. Figure 16 compares their results. We observe that (1) 2EQUAL obtains higher F-measure than 3EQUAL (.64 vs. .54), since it distinguishes between strong and weak indicators for record similarity; (2) LEARNED significantly outperforms the other two strategies (by 50% over 2EQUAL and by 76% over 3EQUAL), showing effectiveness of weight learning. This shows importance of weight learning.

**Attribute contributions:** We then consider the contribution of each attribute for chain classification. Figure 17 shows the results on the perturbed *FBIns* data and we have four observations. (1) Considering only name but not any other attribute obtains a high recall but a very low precision, since all listings on this data set have the same name. (2) Considering dominant-value attributes in addition to name can improve the precision significantly and improve the F-measure by 104%. (3) Considering category in addition does not further improve the results while considering state in addition even drops the precision significantly, since three chains in this data set contain the same wrong value on state. (4) Considering both category and state improves the recall by 46% and




Figure 17: Attribute contribution on perturbed *FBIns*.

obtains the highest F-measure.

**Robustness w.r.t. parameters:** We also ran experiments to test robustness against parameter setting. We observed very similar results when we ranged  $p$  from .8 to 1 and  $\theta_{th}$  from .5 to .7.

### 6.3 Evaluating efficiency

Our algorithm finished in 8.3 hours on the whole *Biz* data set with 18M listings; this is reasonable given that it is an offline process and we used a single machine. Note that simpler methods (we describe shortly) took over 10 hours even for the first stage on fragments of the *Biz* data set. Also note that using the Hadoop infrastructure can reduce execution time for graph construction from 1.9 hours to 37 minutes; we skip the details as it is not the focus of the paper.

**Stage I:** It spent 1.9 hours for graph construction and 2.2 minutes for core generation. To test scalability and understand importance of our choices for core generation, we randomly divided the whole data set into five subsets of the same size; we started with one subset and gradually added more. We compared five core generation methods: NAIVE applies SPLIT on the original graph; INDEX optimizes NAIVE by using an inverted index; SINDEX simplifies the inverted list by Theorem 4.8; UNION in addition merges v-cliques into v-unions by Theorem 4.9; CORE (Algorithm 1) in addition splits the input graph by Theorem 4.10. Figure 18(a) shows the results and we have five observations. (1) NAIVE was very slow. Even though it applies SPLIT rather than finding the max flow for every pair of nodes, so already optimizes by Theorem 4.15, it took 6.8 hours on only 20% data and took more than 10 hours on 40% data. (2) INDEX improved NAIVE by two orders of magnitude just because the index simplifies finding neighborhood v-cliques; however, it still took more than 10 hours on 80% data. (3) SINDEX improved INDEX by 41% on 60% data as it reduces the size of the inverted index by 64%. (4) UNION improved SINDEX by 47% on 60% data; however, it also took more than 10 hours on 80% data. (5) CORE improved UNION significantly; it finished in 2.2 minutes on the whole data set so further reduced execution time by at least three orders of magnitude, showing importance of splitting. Finally, for graph construction, Figure 18(b) shows the linear growth of the execution time.

**Stage II:** After core identification we have .7M cores and 17.3M satellites. It spent 6.4 hours for clustering: 1.7 hours for blocking and 4.7 hours for clustering. The long time for clustering is because of the huge number of blocks. There are 1.4M blocks with multiple elements (a core is counted as one element), with a maximum size of 22.5K and an average of 4.2. On only 35 blocks clustering took more than 1 minute and the maximum is 2.5 minutes, but for 99.6% blocks the size is less than 100 and CLUSTER took less than 60 ms. The average time spent on each block is only 9.6 ms.

### 6.4 Summary and discussions

**Summary:** We summarize our observations as follows.

1. Identifying cores and leveraging evidence learned from the cores is crucial in group linkage.

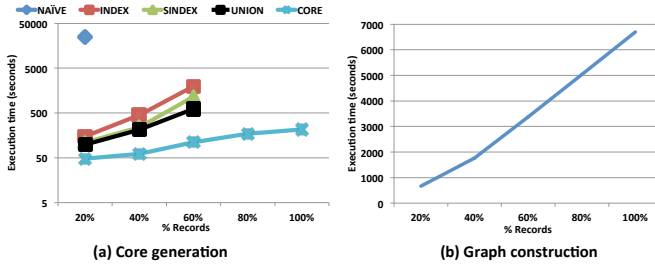


Figure 18: Execution time (we plot only those below 10 hours).

2. There are often erroneous values in real data and it is important to be robust against them; applying ONEDOM and requiring  $k \in [1, 5]$  already performs well on most data sets that have reasonable number of errors.
3. Distinguishing the weights for distinct and non-distinct values, and setting weights of values according to their popularity are critical for obtaining good clustering results.
4. Our algorithm is robust on reasonable parameter settings.
5. Our algorithm is efficient and scalable.

**Discussion:** In the paper, we present single-machine algorithms to identify groups. Performing such data-intensive tasks on powerful distributed hardware and service infrastructures has become popular, in particular with the emerging of widely advisable MapReduce programming model [24, 4, 12]. We next discuss possible parallelized solutions of our algorithms in Hadoop infrastructure.

For graph construction, we can proceed in two steps: (1) to create all cliques where nodes sharing the same common-value and a particular dominant-valued attribute are in the same clique, and (2) to find all maximal cliques. In step (1), we first distribute records and map a record  $r$  to one or more  $\langle \text{key}, \text{value} \rangle$  pairs where  $\text{key}$  is a value on a particular dominant-value attribute of  $r$  and  $\text{value}$  is the value for common-value attribute of  $r$  (Mapper). We then find cliques in each block with a particular  $\text{key}$ , and meanwhile keep an inverted list for each block (Reducer). Step (2) takes the output inverted lists and cliques in Step (1) as input. It first uses each entry in the inverted lists as a  $\langle \text{key}, \text{value} \rangle$  pair to map cliques, so that all cliques that a record  $r$  belongs to are mapped into the same block (Mapper). We then find all maximal cliques within each block (Reducer).

To detect cores in the similarity graphs, the algorithm proceeds iteratively. We can use Spark [31], a cluster computing framework to support iterative jobs while retaining the scalability and fault tolerance of MapReduce. For each iteration, we first partition the input graphs into blocks so that each block contains all records of the same maximal connected component (Mapper), and proceed CORE within each block in parallel (Reducer). Note that the MapReduce solution may not denominate our single-machine solution that takes only 2.2 minutes, because of the additional overhead of the MapReduce program.

In similar ways, we identify groups as follows. We first partition the input elements (satellites and cores) into blocks so that each block contains elements that may potentially belong to the same group (Mapper), and proceed CLUSTER within each block in parallel (Reducer).

## 7. CONCLUSIONS

In this paper we studied how to link records to identify groups. We proposed a two-stage algorithm that is shown to be empirically scalable and accurate over two real-world data sets. Future work includes studying the best way to combine record linkage and group linkage, extending our techniques for finding overlapping groups,

and applying the two-stage framework in other contexts where tolerance to value diversity is critical.

## 8. REFERENCES

- [1] N. Bansal, F. Chiang, N. Koudas, and F. W. Tompa. Seeking stable clusters in the blogosphere. In *VLDB*, pages 806–817, 2007.
- [2] O. Benjelloun, H. Garcia-Molina, D. Menestrina, Q. Su, S. E. Whang, and J. Widom. Swoosh: a generic approach to entity resolution. *VLDB J.*, 18(1):255–276, 2009.
- [3] H. Bruhn, R. Diestel, and M. Stein. Menger’s theorem for infinite graphs with ends. *J. Graph Theory*, 50:199–211, 2005.
- [4] C. Chambers, A. Raniwala, F. Perry, S. Adams, R. R. Henry, R. Bradshaw, and N. Weizenbaum. Flumejava: easy, efficient data-parallel pipelines. In *PLDI*, pages 363–375, 2010.
- [5] W. W. Cohen, P. Ravikumar, and S. E. Fienberg. A comparison of string distance metrics for name-matching tasks. In *IWEB*, 2003.
- [6] D. Dey. Entity matching in heterogeneous databases: A logistic regression approach. *Decis. Support Syst.*, 44:740–747, 2008.
- [7] X. Dong, A. Y. Halevy, and J. Madhavan. Reference reconciliation in complex information spaces. In *Proc. of SIGMOD*, 2005.
- [8] A. K. Elmagarmid, P. G. Ipeirotis, and V. S. Verykios. Duplicate record detection: A survey. *IEEE Trans. Knowl. Data Eng.*, 19(1):1–16, 2007.
- [9] S. Even and E. R. Tarjan. Network flow and testing graph connectivity. *SIAM Journal on Computing*, 4(4):507–518, 1975.
- [10] W. Fan, X. Jia, J. Li, and S. Ma. Reasoning about record matching rules. *PVLDB*, 2(1):407–418, 2009.
- [11] I. P. Fellegi and A. B. Sunter. A theory for record linkage. *Journal of the American Statistical Association*, 64(328):1183–1210, 1969.
- [12] R. L. Ferreira Cordeiro, C. Traina, Junior, A. J. Machado Traina, J. López, U. Kang, and C. Faloutsos. Clustering very large multi-dimensional datasets with mapreduce. In *KDD*, pages 690–698, 2011.
- [13] L. R. Ford and D. R. Fulkerson. *Flows in networks*. Princeton University Press, 1962.
- [14] T. Gonzalez. On the computational complexity of clustering and related problems. *Lecture Notes in Control and Information Sciences*, pages 174–182, 1982.
- [15] S. Guo, X. Dong, D. Srivastava, and R. Zajac. Record linkage with uniqueness constraints and erroneous values. *PVLDB*, 3(1), 2010.
- [16] O. Hassanzadeh, F. Chiang, H. C. Lee, and R. J. Miller. Framework for evaluating clustering algorithms in duplicate detection. *PVLDB*, pages 1282–1293, 2009.
- [17] M. A. Hernandez and S. J. Stolfo. Real-world data is dirty: Data cleansing and the merge/purge problem. *Data Mining and Knowledge Discovery*, 2:9–37, 1998.
- [18] S. Huang. Mixed group discovery: Incorporating group linkage with alternatively consistent social network analysis. *International Conference on Semantic Computing*, 0:369–376, 2010.
- [19] N. Koudas, S. Sarawagi, and D. Srivastava. Record linkage: similarity measures and algorithms. In *SIGMOD*, 2006.

- [20] B. Larsen and C. Aone. Fast and effective text mining using linear-time document clustering. In *KDD*, pages 16–22, 1999.
- [21] P. Li, X. L. Dong, A. Maurino, and D. Srivastava. Linking temporal records. *PVLDB*, 4(11):956–967, 2011.
- [22] X. Liu, Y. Gong, W. Xu, and S. Zhu. Document clustering with cluster refinement and model selection capabilities. In *SIGIR*, 2002.
- [23] B. W. On, N. Koudas, D. Lee, and D. Srivastava. Group linkage. In *ICDE*, pages 496–505, 2007.
- [24] K. Shvachko, H. Kuang, S. Radia, and R. Chansler. The hadoop distributed file system. In *MSST*, pages 1–10, 2010.
- [25] P. R. Silhouettes. A graphical aid to the interpretation and validation of cluster analysis. *Journal of Comp. and Applied Math.*, 20(1):53–65, 1987.
- [26] J. Sima and S. E. Schaeffer. On the np-completeness of some graph cluster measures. *Lecture Notes in Computer Science*, 2006.
- [27] S. E. Whang, D. Menestrina, G. Koutrika, M. Theobald, and H. Garcia-Molina. Entity resolution with iterative blocking. In *SIGMOD*, 2009.
- [28] D. T. Wijaya and S. Bressan. Ricochet: A family of unconstrained algorithms for graph clustering. In *DASFAA*, pages 153–167, 2009.
- [29] W. E. Winkler. Methods for record linkage and bayesian networks. Technical report, U.S. Bureau of the Census, 2002.
- [30] M. Yoshida, M. Ikeda, S. Ono, I. Sato, and H. Nakagawa. Person name disambiguation by bootstrapping. In *SIGMIR*, 2010.
- [31] M. Zaharia, M. Chowdhury, M. J. Franklin, S. Shenker, and I. Stoica. Spark: cluster computing with working sets. In *HotCloud*, 2010.